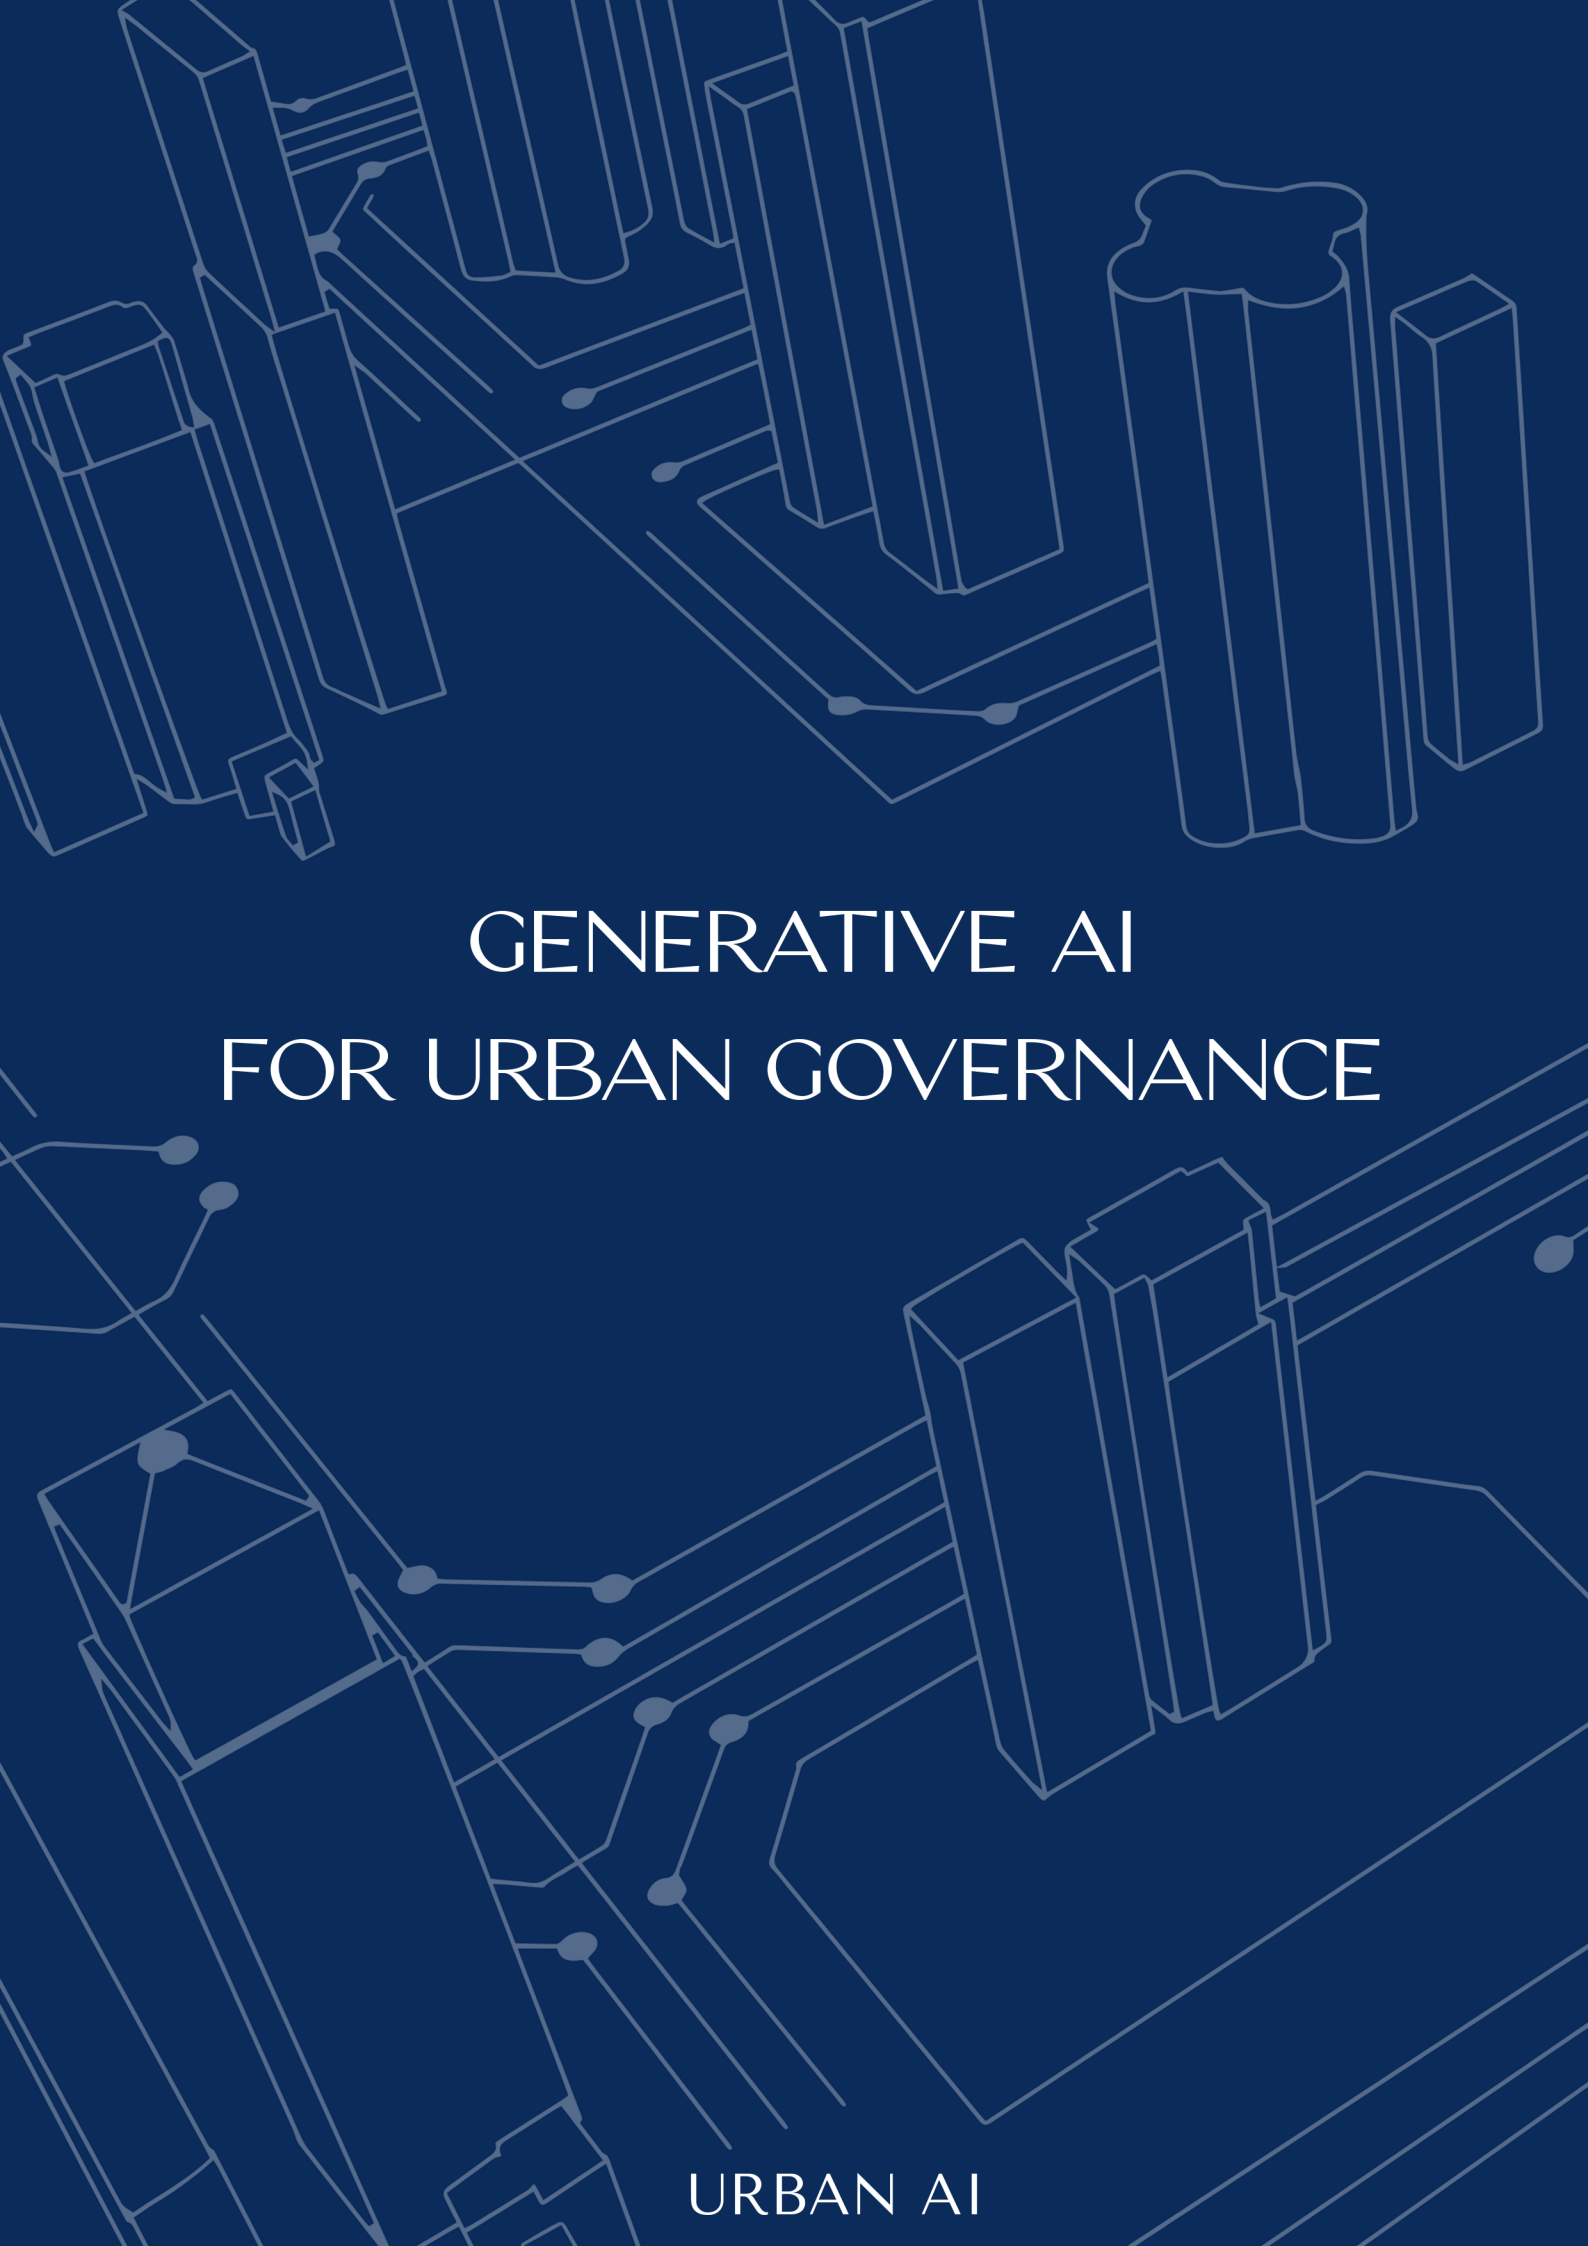

GENERATIVE AI FOR URBAN GOVERNANCE

URBAN AI

ACKNOWLEDGMENTS

We would like to thank the experts, practitioners, and public officials who provided insights from which the recommendations of this report are based. More broadly, we would like to express our gratitude to the entire Urban AI Community as well as Urban AI Partners for their continuous support during this whole project. Without their contribution, this report would not be possible.

SUGGESTED CITATION

Nelson, RJ. Naing, KM. Beroche, H. 2023. Generative AI for Urban Governance. Urban AI

ABOUT THE AUTHORS

Ruth Nelson

Ruth Nelson is an architect and spatial data scientist currently working as an interdisciplinary PhD researcher in the Centre for Urban Science and Policy at Delft University of Technology in the Netherlands, with both a Master's in Architecture (Nelson Mandela University) and MRes. in Space Syntax (University College London). Ruth was selected as an Emerging Leader in Urban AI and is now a member of Urban AI. She has worked in corporate, research and private consulting roles, across South Africa, the UK, Mexico and the Netherlands. Her research and practice is focused in supporting the development of more equitable policies and design solutions through harnessing the power of digital technologies and spatial data science. For further information, refer to: https://www.ruthjnelson.com/.

Khant Min Naing

Khant Min Naing is an undergraduate student at Singapore Management University, double majoring in Smart-City Management Technology and Public Policy. He has garnered over 5 years of work experience across government offices, non-governmental organisations, and research think-tanks in Myanmar, Singapore, Hong Kong, and France. His research areas include geospatial data science and analytics, computational social science, and urban policy and planning. He is particularly interested in the challenges and opportunities of urbanisation in Southeast Asian cities, and is committed to contributing to the academic research on the region. Currently, Khant is interning as a junior researcher at Urban AI.

Hubert Beroche

As the founder of Urban AI, Hubert institutionalised the study of urban artificial intelligence, and has built a large, global network of technologists, municipalities, start-ups, researchers, and other urban practitioners who interact with artificial intelligence in urban contexts. He collaborated with several worldwide cities and urban stakeholders to design and implement Urban Artificial Intelligence. He is the author of several pioneering works in this field. He is also a Lecturer at Sorbonne Université and has a background in Urban Economics, Climate Policy and Management.

TABLE OF CONTENT

1. INTRODUCTION

According to the United Nations approximately 56% of the world's population live in cities, with this number set to reach nearly 70% by 2050. Cities form the heart of human development and innovation, with 80% of global GDP generated within them (World Bank, 2023). Since the launch of ChatGPT at the end of 2022, the world has experienced a proliferation of debates centred on generative AI, particularly in relation to the opportunities it will provide for both personal and collective productivity. A recent report by Mckinsey (2023) notes that generative AI's impact on productivity could add trillions of dollars in value to the global economy. However the speed with which this technology is being adopted is also leading to concerns over potential security threats and biassed outcomes (Baxter & Schlesinger, 2023).

Less attention has been focused on the impact that this technology might have for urban governance, compared to other industries. This may be due to governments generally being slower to adopt technology than the private sector as a result of a number of factors, such as lack of funding, higher public scrutiny, complex contracting processes, lack of internal IT capacity, and agency fragmentation (Hinkley, 2023). However, over the last few decades, increasingly cities are adopting digital technologies and the use of AI within their management processes, such as the exploration of the use of digital twins in many countries like the UK, the development of smart cities such as Sejong in South Korea, deployment of crowd sensing technology in Amsterdam and the release of public data by cities such as New York. It is thus conceivable, in this digital age, that generative AI may be integrated into processes within city governments.

We are only beginning to understand generative AI's capabilities and risks. This report represents an initial effort to further our understanding of generative AI's potential impact on urban governance. We conduct a literature review of existing research to understand how generative AI might be applied in the context of cities, followed by analyses of 10 semi-structured interviews with experts across industry, government and academia.

We kindly thank and had the pleasure of interviewing:

- Anthony Townsend — Urbanist in Residence, Cornell Tech
- Ariel Noyman — Urban Scientist, MIT Media Lab
- Shazade Jameson — Senior Consultant Tech Governance & Urban AI
- Maria Mamoura — Director, Bryden Wood
- Alexander Kamenev — Founder, Aino.World
- Paula Boet Serrano — Project Manager, Barcelona City Council

• Zhongwen Huang	Director - Smart City Projects Office, Smart Nation and Digital Government Office Singapore
• Ernest Kwan	Assistant Deputy Minister, Department of Families, Government of Manitoba
• Emily Binet Royall	Smart Cities Administrator, City of San Antonio
• Santiago Garces	Chief Information Officer, City of Boston

2. LITERATURE REVIEW

The advent of big data and increases in computing power is allowing for the progressive adoption of Artificial Intelligence (AI) characterising a new era of urban research, planning and policy (Chui et al., 2022; Kandt and Batty, 2021). Artificial intelligence was originally coined by the Stanford Professor, John McCarthy in 1955 as "the science and engineering of making intelligent machines". More recently, Russell and Norvig (2010) define AI as the "designing and building of intelligent agents that receive precepts from the environment and take actions that affect that environment". Whereas the High-Level Expert Group on Artificial Intelligence (AI HLEG) of the European Commission (2018) goes one step further stating that AI are not merely agents, but systems that display intelligent behaviour by analysing their environment, taking actions and to a certain extent, autonomously achieving specific goals.

As AI technology continues to evolve, the autonomous agency of AI and its ability to immediately or eventually replace human intelligence is a hot topic of debate. Researchers predict that AI will surpass human performance in many tasks in the next few decades, including "driving trucks (by 2027), working in retail (by 2031), writing bestsellers (by 2049), and performing surgery (by 2053)" (Grace et al., 2018:3). They also believe that AI could outperform "all human tasks within 45 years", with all human jobs eventually replaced "within 122 years" (Grace et al., 2018:3). Alternatively, others advocate that human intelligence will not fundamentally be replaced, but rather augmented and enhanced by AI (Cremer and Kasparov, 2021; Kimm, 2022). The increasing adoption and autonomy of AI is also leading to many ethical considerations, relating to bias, privacy and safety concerns (Jaupi, 2022; O'Sullivan, 2021).

The recent surge in media attention garnered by many Generative AI models and services, such as DALL-E, Midjourney, Bard and ChatGPT, has only served to heighten this debate. Whilst these specific models are gaining widespread popularity, less is spoken about many other kinds of generative AI models which exist and are being employed by practitioners and researchers for a range of applications and task automation. **Thus, the purpose of this review is to shine light on different kinds of existing generative AI models and the types of problems they are currently being applied to, before presenting our research on the potential of generative AI to assist in decision making processes in the management of cities.**

2.1. The Architecture of Generative AI

Generative AI can be understood as a subset of artificial intelligence, with the special capability of generating new and realistic representations of various

forms of data, including not limited to text, images, and sound (Gioti, 2021; Huzaifah & Wyse, 2021). Compared to conventional machine learning and deep learning models, generative AI models frequently have more complicated architecture and development processes, making them a more specialised subset of AI models (refer to Figure 1.1). Hence, they are able to perform more than just learning and predicting patterns and structures from the given dataset (which machine learning and deep learning models are typically tasked with) by generating new data that resembles the training data set.

For example, a traditional machine learning model and a generative model can be trained on the same dataset of historical street view images of a city. The machine learning model would learn to identify and classify different features of the images, such as roads, buildings, and vegetation, and can predict the liveability or walkability of the city. However, the generative model, whose training and development is more computationally-sophisticated, can not only predict but also generate new street view images of the city based on different scenarios, such as if a new park is built or if the road structure is changed. In this way, generative models provide richer insights and information and can revolutionise the way AI models are used.

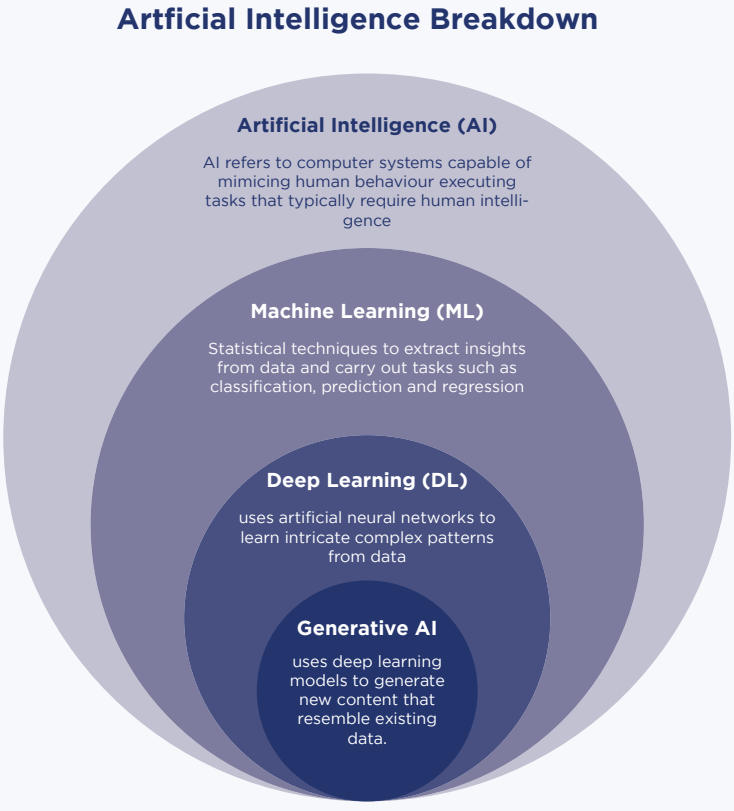

Figure (1.1): Breakdown of Artificial Intelligence and its wide range of models with diverse capabilities and applications.
Reference: Purohit (2023)

The discourse around generative AI has been mainstreamed thanks to the recent advancement of two classes of deep learning models in particular: the Generative Adversarial Networks (GAN) and Generative Pre-Trained Transformer (GPT). Both GANs and GPTs are used for generating new forms of data, but they have different architectures, parameters and employ different algorithms.

GANs define a specific learning procedure that involves training two neural network block[1] (a stack of multiple neural network algorithms) - a generator and a discriminator. The generator network generates synthetic data samples that are similar to a given training dataset, and then training a discriminator network to distinguish between real and synthetic samples, refer to Figure 1.2. The goal of this process is to find a generator that can produce highly realistic synthetic data samples that are similar to the original training data. Once a GAN model has been trained on a dataset, it can be used to generate new data samples that are similar to the training data, such as new images or other types of datasets such as population data. A simplified architecture of a GAN model is illustrated in the figure below.

Generative Adversarial Network Architecture

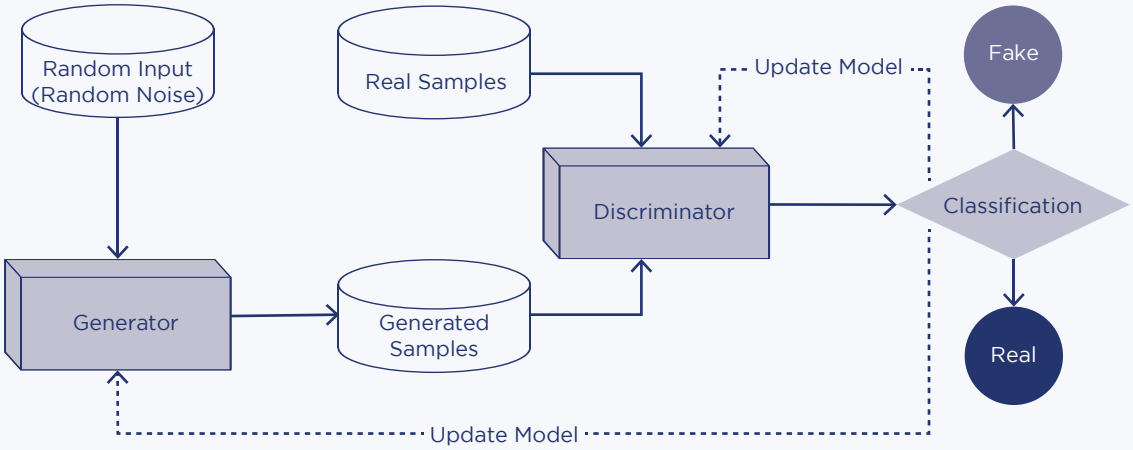

Figure (1.2): Architecture of Generative Adversarial Network (GAN) Model
Reference: By Authors

GPT models use a stack of Transformer (Vaswani et al., 2017) layers, each of which has two main components: self-attention and feed-forward networks. Self-attention allows the model to learn long-range dependencies in the input sequence, while the feed-forward networks allow the model to learn complex non-linear relationships between the words. Together, this process of deep learning trains the model to understand the intricacies of language, allowing it to predict the likelihood of a sequence of words given a context. Once a GPT model has been trained, it can be used to generate new text that is realistic and similar to the input data. Due to this nature, GPTs are mostly used for text-producing tasks, such as text translation, text summarization and question answering. The earliest pre-trained language models include Word2vet (Mikolov et al., 2013) and BERT (Devlin et al., 2018). Subsequent models built on the BERT architecture, including GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020). Both DALL·E (text-to-image) and ChatGPT (text-to-text) are built on the GPT-3 architecture. A simplified architecture of a GPT model is illustrated in Figure 1.3.

[1]Artificial neural networks (shortened to neural networks or neural nets) are a branch of machine learning models in which a computer learns to perform some task by analysing training examples. Such models are inspired by biological neural networks that constitute the animal brain. They “learn” to perform tasks by examples rather than any task-specific rules.

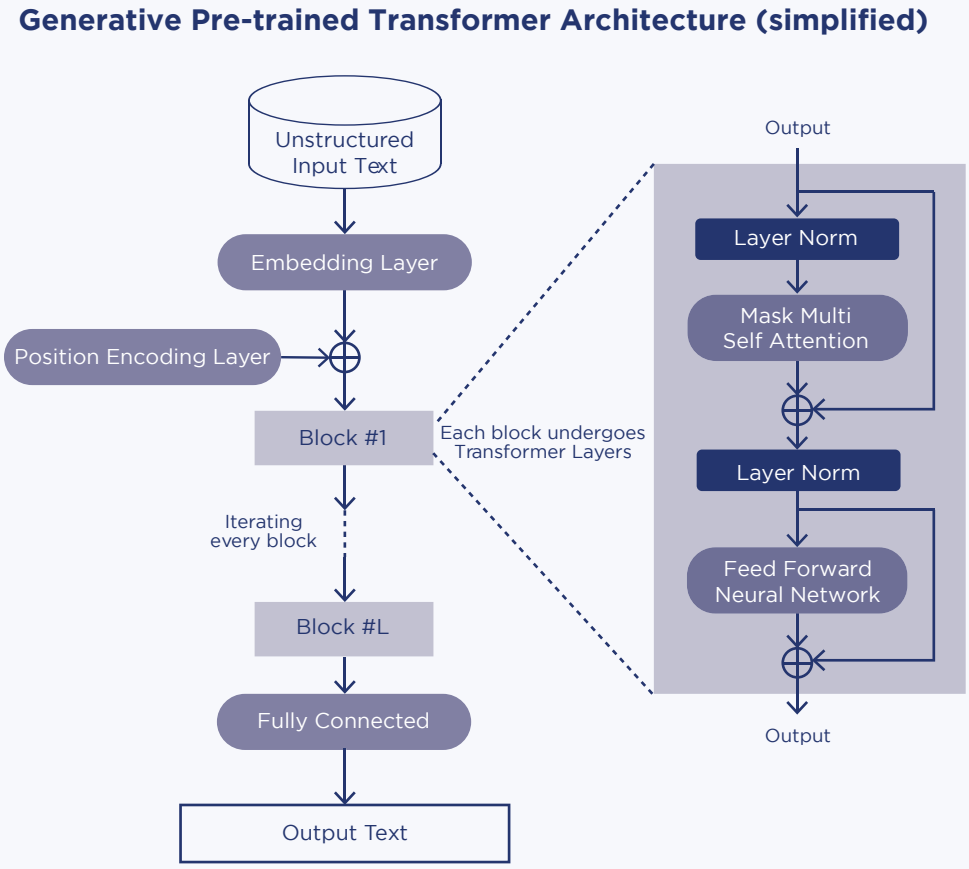

Figure (1.3): Architecture of Generative Pre-trained Transformer (GPT) Model

Reference: by Authors

GPTs and GANs are both generative models, but they have different use cases. GPTs are better at understanding and generating sequential data such as language, while GANs are better at generating other forms of data, such as images, which has been more widely explored within applications for cities.

GPTs and GANs are both generative models, but they have different use cases. GPT models are based on the transformer architecture, which is well-suited for processing sequential data, such as text and code. This allows GPT models to understand the dependencies between words and phrases, and generate text that is both grammatically correct and semantically meaningful. On the other hand, GAN models undergo the adversarial training process, which continuously train itself to enhance and maximise the resemblance of the output data in comparison to the training data. This makes them well-suited to tasks such as synthetic data generation or visual-aided simulation, which has been more widely explored within applications for cities.

2.2. Generative AI in Cities

Cities are being conceived as key sites in the production of data (Kitchin, 2015). The term ‘big data’ refers to massive, dynamic, varied and detailed datasets that can be connected and utilised in diverse ways (Kitchin, 2013). The rise of the term, the “real-time” city has emerged as big data is increasingly collected and analysed in real time by cities for various applications such as automatic

traffic fines (Nizzad et al., 2021; El Hansali et al., 2022; Mandal et al., 2020), the deployment of police and other public service vehicles to certain areas (Ellison et al., 2021) and the monitoring of environmental conditions, such as air quality (Kang et al., 2018).

In response to this deluge of data researchers are exploring the use of Generative AI applications within urban settings, particularly GAN-based models. The high-dimensionality of these generative models allow them to generate complex and detailed outputs that can be useful in the built environment and urban planning. Outputs can be varied at different scales and formats based on the type of input data such as satellite imagery, population data, traffic patterns, and environmental data (Wang et al., 2022). Wu et al. (2022) introduce a framework that categorises different generative models currently being researched or experimented in for application in urban settings based on the type of generated output - images, structured data and graphs, refer to Figure 2.1. By clearly defining the different types of generative models, this framework can help researchers better understand the strengths and limitations of each model and make more informed decisions about which approach to use in their work.

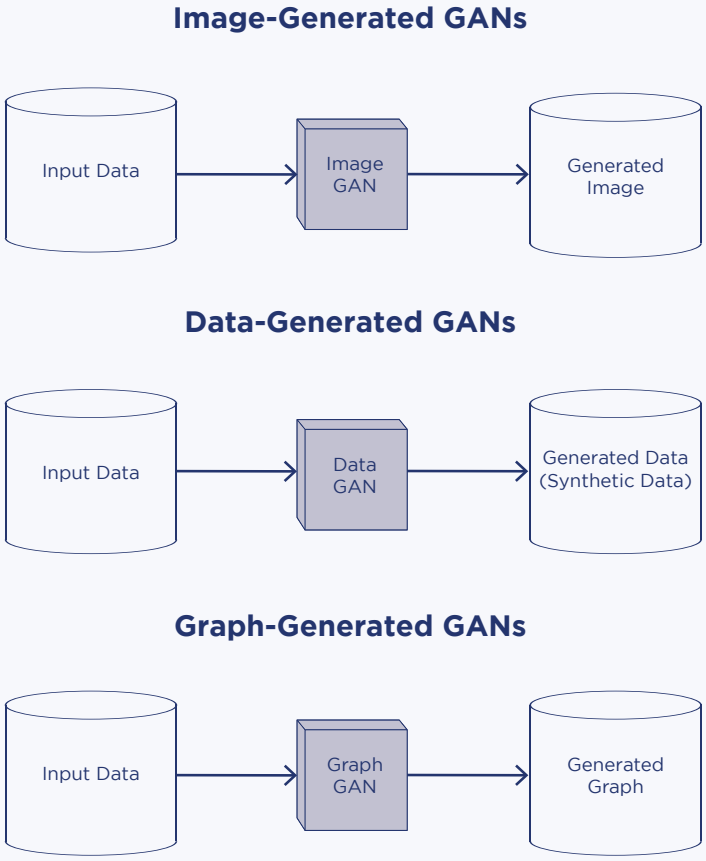

Figure (2.1): Classification of GANs based on the Type of Generated Outputs
Reference: Wu et al. (2022)

Image-generated models produce a wide variety of images such as landscapes, street views, maps and satellite images with hyper realistic details. Examples of application include using images of existing land use designs to create new areas such as Pix2Pix a model used to create novel urban impressions iteratively and rapidly (Raman, Kollar, and Penman, 2022). The generated urban impressions provide alternative land use formations for planners and the municipalities to critique and discuss. Alternatively in a different study, GANs are used to generate terrain maps with different layers of building footprint, road networks, vegetation and water body with the input of semantic masks of street views (Kim et al., 2019). GANs are reported to be used in transformation of input satellite images into new map images that are rendered in a style that is consistent with a reference map (Xu and Zhao, 2018; Kang et al.,2019; Zhang et al., 2020; Li et al., 2020). Additionally, in more sophisticated and automated applications, high-functional GANs have been demonstrated to generate fully vectorized levels of building models. Such models are derived from readily available, below-resolution Digital Surface Models (DSMs), which are generated from stereo satellite images. Wang et al. (2021) provide evidence of this capability of GANs in their study.

Automated Vectorisation of Satellite Images

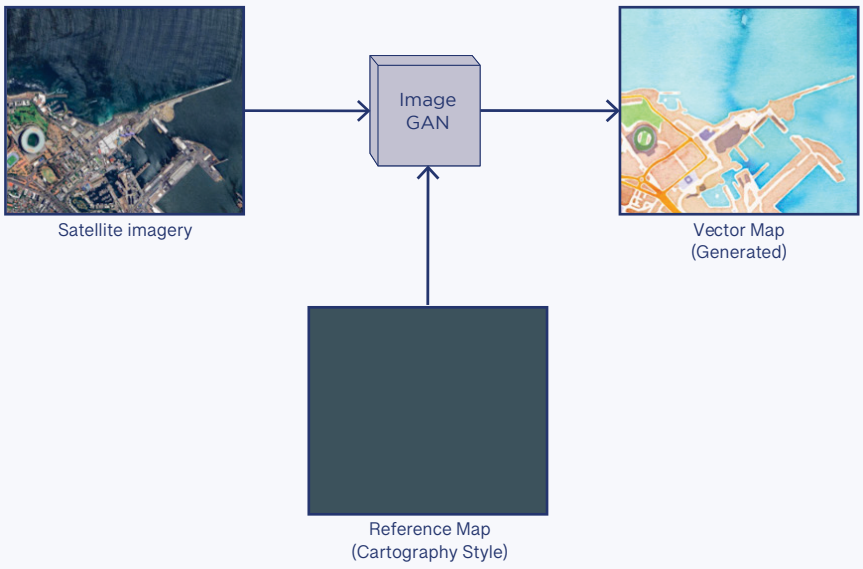

Figure (2.2): An example use-case of Image-generated GAN for automated vectorisation - the model takes an input image and generates a corresponding vectorized image, which can be used for various applications. The visuals used in the figure are not the actual output of the model, but rather an abstract representation of how it works. The accuracy of the actual model output may vary.

Reference: by Authors

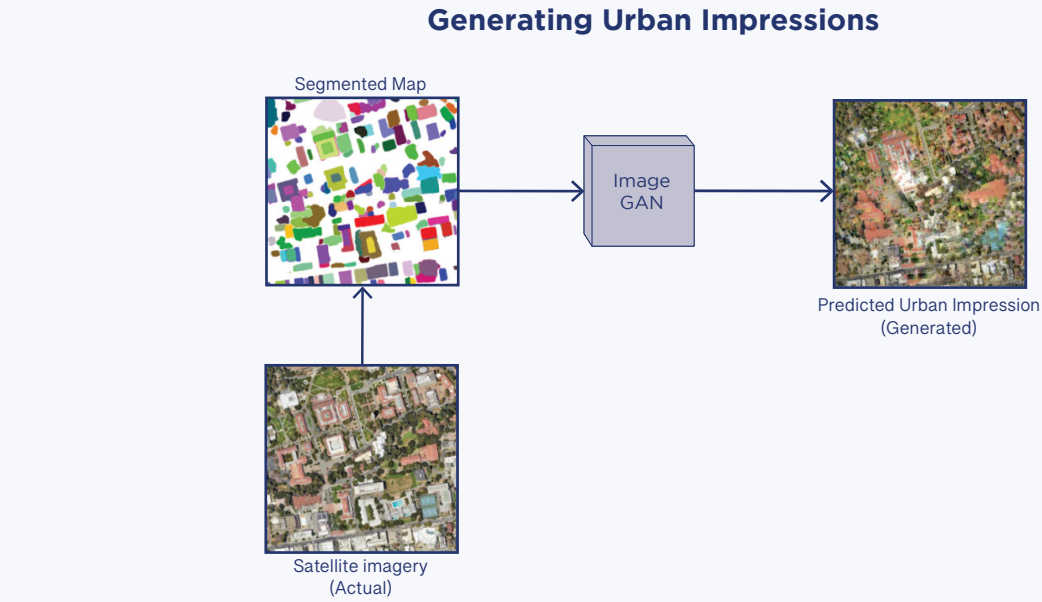

Figure (2.3): An example use-case of Image-generated GAN for generating urban impressions - the model takes an segmented image derived from an actual satellite image and generates a predicted urban impression, which resembles the actual image. The visuals used in the figure are not the actual output of the model, but rather an abstract representation of how it works. The accuracy of the actual model output may vary.

Reference: by Authors

Graph-generated models process graph data networks as input or output and enable the description of spatial relationships at different scales. For example, they can be used to generate variants for the configuration of urban design elements, such as street networks and transportation networks from geographical input data (Bielika et al., 2019) or the context embedding of a virgin area (Wang et al., 2022).

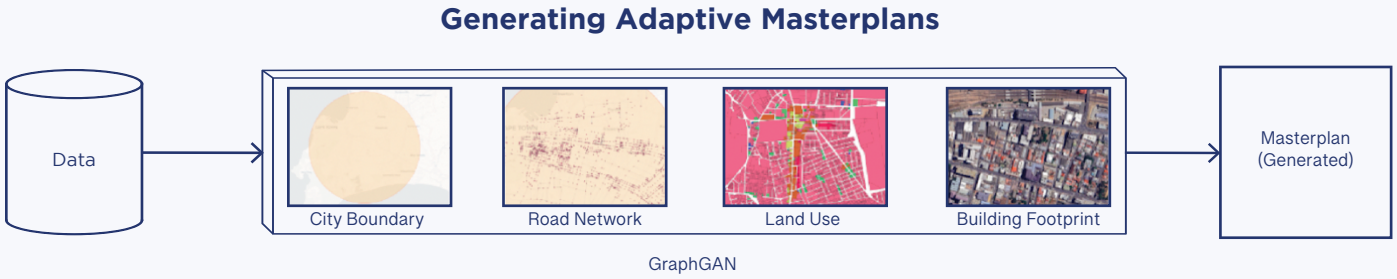

Figure (2.4): An example use-case of Graph-generated GAN for generating adaptive masterplan - the model takes an array of geographical input data and generates various configurations of urban design elements (e.g., streets, plots, buildings, land use). The visuals used in the figure are not the actual output of the model, but rather a representation of how it works. The accuracy of the actual model output may vary.

Reference: Bielika et al. (2019)

Data-generated models process vast amounts of data such as text, GPS, and mobility data to generate new data samples. These models are particularly useful in cases where the actual data is not available or data collection is challenging and constrained. By generating new samples that are similar to the "ground truth," or the actual data, these GANs can improve downstream

learning of the model. For example, GANs can be used to simulate realistic urbanisation patterns under various conditions by studying the complex spatial organisation observed in global urban patterns (Albert, 2018). GANs can also be utilised to generate realistic traffic conditions based on not-yet-observed travel demands by learning the fundamental patterns of how traffic conditions evolve with changes in travel demand and the underlying structure of the road network (Zhang, 2020). Additionally, GANs can be utilised to estimate future housing growth patterns from historic satellite maps (Ilbrabim et al., 2021) and to create synthetic demographic profiles, including the type and composition of households, income, and social demographics, for urban residents in new neighbourhood developments (Johnsen, 2020).

As discussed above, researchers have explored the use of generative AI models for a wide range of applications in urban domains. Table 1 has provided a summary on the different application areas of all three types of GAN models for easier reference.

Table 1: Summary of GAN Models and Application Areas

Application	Type	Tool / Model	Description of Application
Creation of Urban Impressions (Raman, Kollar, and Penman, 2022)	Image	Pix2Pix	To create novel urban impressions[2] to be used at the macro-level as an iterative and rapid method for imagining, critiquing and discussing alternative urban formation across whole neighbourhoods.
3D Urban Construction (Kim et al., 2019)	Image	DCGANs	To generate terrain maps (in different levels of building, road, vegetation, water body etc.) from semantic masks of street views.
Cartographic Style Transfer (Xu and Zhao, 2018; Kang et al.,2019; Zhang et al., 2020; Li et al., 2020)	Image	Pix2Pix CycleGAN SG-GAN MapGAN	To input a content image and a style reference image to generate a map that retains the information of the content map but painted in the style of reference map.
Vectorization of Satellite Images (Wang et al. 2021)	Image	cGAN	To generate fully vectorised levels of building models from readily available below resolution DSMs generated from stereo satellite image.

[2]Urban impressions refer to a set of procedurally generated, aerial-style images of an urban area. Urban impressions are not intended to recreate perfectly realistic aerial imagery of an area; rather, they are focused on creating a perception of certain urban characteristics.

Application	Type	Tool / Model	Description of Application
Land Use Configuration (Wang et al. 2022)	Graph	LUCGAN	To generate the land-use solution (that is tailored to community needs) based on the context embedding of a virgin area.
Adaptive Master Plans (Bielika et al., 2019)	Graph	AMP	To generate variants for the configuration of urban design elements (e.g., streets, plots, buildings, land use) from geographical input data.
Estimating Traffic Conditions (Zhang, 2020)	Data	TrafficGAN	To generate realistic traffic conditions given a not-yet-observed travel demand by learning the fundamental patterns of how traffic conditions evolve with changes in travel demand and the underlying structure of the road network.
Population Synthesis (Johnsen, 2020)	Data	CVAE & cGAN	To generate a synthesis demographic profile (e.g., type and composition of households, income, and social demographics) of urban residents in new neighbourhood developments.
Housing Pattern Simulation (Ilbrabim et al., 2021)	Data	GAN	To estimate future housing growth patterns from historic satellite maps
Urbanisation Simulation (Johnsen, 2020)	Data	GAN	To stimulate realistic urbanisation patterns under different conditions in cities by learning complex spatial organisation observed in global urban patterns.
Land Use & Land Cover Prediction (Sun et al., 2021)	Data	Pix2Pix	To predict and generate the future land use and land cover change by learning historic urbanisation patterns.

2.3. Generative AI and Urban Governance

With improved availability of big data for urban governance, the potential of generative AI for the management of cities also grows. By effectively harnessing big data, it is conceivable that generative AI could empower alternative visions for city development, which could be integrated into existing data and analytics programs within city governance. During the past few decades, governance has rapidly evolved with the integration of big data analytics (Kitchin and Lauriault, 2014) and many cities have embraced the importance of data availability. For example, New York City launched its Open Data Portal, which makes the public data generated by various city agencies and other organisations available for public use. Another instance can be found in Amsterdam, where the government launched the Amsterdam Smart City initiative, which leverages public-private partnerships and data analytics to drive multiple smart city projects (Fitzgerald, 2016). Similarly, the Helsinki Region Infoshare (HRI) platform is another notable example of municipal data platform and systematically releases open city datasets and interfaces for public use (Hämäläinen, 2020).

Many questions remain if and whether the application of generative AI could potentially offer more efficient and effective management processes than current practices. It is a crucial question that needs to be explored further, given the risks, ethical concerns, and explainability challenges that are often associated with the use of AI technology. Raman et al (2022) note that the utility of these AI tools in practice can still be very challenging for many planning and design related tasks, without contracting out or hiring in specialised skill sets. Moreover, the decisions made and generated by AI systems may not always be transparent, accountable, or explainable, leading to the “black box” problem and a need for explainability of the AI decision-making (Marshan, 2021).

In 2020, European Parliamentary Research Service (EPRS) released a study which discussed the ethical and moral risks associated with the deployment of AI technologies, as well as initiatives that exist to address them (Bird et al., 2020). In terms of legal framework, the recent AI Act has received the approval of the key parliamentary committees in the European Parliament in a vote held in May of 2023 (Bertuzzi, 2023). The Act outlines a set of rules for the development and application of AI, including a list of prohibited AI practices, mandatory testing and documentation requirements, and penalties for non-compliance. This regulation has been subject to debate and criticism from various stakeholders, with some arguing that it could stifle innovation and leveraging the potential of AI utility (Bertuzzi, 2023).

In different parts of the world, cities are beginning to harness the power of data analytics, and high level government legislation is being developed in relation to the application of AI. However, there is limited knowledge on the interface between local municipalities and AI, particularly generative AI, as well as the role of local governments in reaping and leveraging the potential applications and managing related risks of this technology. This leads to the very purpose of this report, to explore these inquiries through interviews with city stakeholders and experts from various parts of the world.

Disclaimer

The insights presented in the subsequent sections are the synthesis of the interviews with the expert panel and hence are a reflection of their current knowledge and expertise up to the present time. Urban AI makes no guarantee, either expressed or implied, that the challenges and opportunities explored in this research are complete and will fully materialise. This research is intended to serve only as a snapshot of current landscape and debates on the technology and Urban AI would like to call for further exploration and discussion or challenge against.

3. FINDINGS

To address the gaps of knowledge in terms of the interface between municipalities and generative AI, we draw on an assessment of the potentials and pitfalls that generative AI holds for city municipalities. We used a multi-stakeholder approach by conducting a total of 10 interviews with our expert panel, which comprise city stakeholders from various domains such as technology, engineering, academic research and municipal governance. Based on these interviews, we identified and synthesised multiple hypotheses, perspectives and considerations regarding the application of generative AI for the public sector. These insights were also backed and supported by emerging trends and example use cases of generative AI globally. Together, the findings will be presented in four parts as follows.

- **Section 3.1 : Applying Generative AI within City Governance**, which explores the existing and potential domains and applications of Generative AI in terms of city governance.
- **Section 3.2 : Role of Municipalities in Developing Generative AI**, which looks at the roles that municipalities can play in developing and integrating Generative AI tools for government workflows or public use cases.
- **Section 3.3 : Skills and Workflows**, which discusses aspects of the skill, capacity and procedural requirements for successfully developing, integrating and governing generative AI tools.
- **Section 3.4 : Challenges and Future of Generative AI** which reflects on the future dynamics of generative AI and outlines the challenges and considerations to be made.

In combination, this section provides opportunities for city stakeholders to anticipate the possibilities of generative AI and critically assess and mitigate the risks and vulnerabilities associated with the rapid evolution of this technology.

3.1. Applying Generative AI within City Governance

Recent developments in the user interface of generative AI, such as have been observed with ChatGPT or Midjourney, has allowed for greater accessibility of this technology by the general public. **Ariel Noyman** (Urban Scientist, MIT Media Lab) notes in his interview that many AI models have enhanced their user interfaces to enable almost anybody with an internet connection, regardless of their technical proficiency, to interact with them. This exposes many city governments, who might not possess the internal technical competencies, to leverage this technology in-house. The question that frequently arises in discussions about AI, and in this case generative AI, is how will this emerging technology influence and transform existing processes and workflows within governments, which we address in this section.

3.1.1. Productivity Enhancement

> Even before the current craze of generative AI, our approach to tool making is always that we are making tools.
>
> We are not replacing humans in the process. In fact, the whole idea of this kind of work is around augmenting people, creating what we call augmented intelligence.
>
> *Ariel Noyman*
> *Urban Scientist, MIT Media Lab*

Generative AI holds potential to serve as a human "co-pilot" within government operations, serving as a productivity enhancer within the existing workforce. **Santiago Garces** (Chief Information Officer, City of Boston) reflected on how the City of Boston has recently adopted the use of enterprise grade Bard for all government employees and how this is augmenting government workflows, both for core administrative tasks, such as drafting job descriptions, generating alternative policy viewpoints, as well as aiding in routine tasks such as writing meeting memos, and drafting press releases. This is a real world example which shows how generative AI can improve productivity and free up more time and resources for government officers to focus on other complex tasks that perhaps require greater human judgement and reasoning. Whilst the growing concern of eventual job replacement by AI is an important consideration for the long term future that governments have to aware of and prepare for, there was a consensus amongst the interviewees that government employees currently tend to be very stretched, and thus the use of generative AI could serve as an opportunity to actually enhance their productivity, improve working conditions and in doing so provide better services for citizens.

3.1.2. Planning and Management of the Built Environment

A strong theme that emerged across the interviews is that generative AI could play a role in expediting decision-making processes in relation to the planning and management of the built environment. There is potential for municipalities to harness the computational capabilities of generative AI to generate a diverse number of design alternatives which satisfy a complex plethora of demands and considerations (Ochoa, 2023) such as constraints imposed by different regulations relating to infrastructure, road network and dynamic factors such as population growth. Furthermore, what sets generative AI apart from conventional AI models is its capability to generate recommendations and outputs in a multitude of data formats.

According to **Ariel Noyman**, generative AI tools have the ability to create and craft various outputs, such as text, images, videos and three-dimensional models, which were previously complex, time-consuming and costly tasks. Municipalities, together with designers and planners, could adopt these tools to facilitate the exploration of different scenarios and alternatives in city development. For instance, AI-generated renderings such as those produced by UrbanistAI can be utilised to evaluate whether a building meets the necessary regulations and planning requirements. This approach offers a potentially more efficient way for municipalities to assess future developments and encourages dialogue regarding the advantages and disadvantages of proposed changes. Additionally, **Ariel Noyman** mentions the potential of using generative AI to generate physical and sensory experiences of planned development through integrating with augmented reality (AR) and virtual reality (VR) technologies.

Alexander Kamenev (Founder, Aino.World), in his interview, suggests that the insights gained from generative AI may assist in informing decision-making regarding the design and management of urban spaces, allowing municipalities to minimise risks. The use of generative AI allows for opportunities to streamline risk management through simulating different outcomes, which could be used to explore how different factors and events could impact the physical environment. For example, by analysing spatialised data such as historical flooding records or changes in land use, municipalities and planners could generate scenarios that depict how different physical spaces might be affected by future events like floods. This capability would be particularly significant for countries like India, Bangladesh and Vietnam, where the risk of flooding is a pressing concern. With this kind of potential technology at disposal, municipalities could take a proactive approach in urban planning and design, not merely through images or text, but through inputting spatial data to generative models.

3.1.3. Augmenting Public Services

Data-driven Insights

The use of generative AI models to gain insights from unstructured data in the city context could be a valuable approach. Municipalities face notable challenges in analysing datasets gleaned from the city for management purposes, which is often unstructured, such as images, video footage, recordings and textual data. Conventional approaches to data analysis often require a significant investment of time and effort in manually structuring the data before conducting any analysis. However, generative models can assist in clustering, classifying, and analysing unstructured data to **re-produce** the data into structured and organised datasets at minimal time and labour cost. In this regard, municipalities can leverage generative AI models to rapidly process and analyse data from various sources. **Ariel Noyman** refers to this as “the real-time city” where municipalities collect data through sensors in real-time and achieve a close-to-reality understanding of what is currently happening in the city - from movement of behaviour to changes in the physical environment. It is only made possible through the effective leverage of data and technological

tools. Municipalities have started using AI to accumulate huge amounts of urban data - for example, the city of Rio de Janeiro in Brazil and New York city have sought to build integrated data analytics centres that weave together data streams from a diverse set of city agencies and sensor networks. Through such efforts, generative AI could significantly reduce the processing time and cost of unstructured data to allow for improved real time management of cities.

Public engagement

> Obviously a city like our city is always constantly trying to find new ways to engage our citizens and collect that feedback to inform policies, initiatives and projects.
>
> *Ernest Kwan,*
> *Assistant Deputy Minister, Department of Families,*
> *Government of Manitoba*

Cities are constantly seeking new ways to engage citizens and collect feedback to inform policies and initiatives. **Ernest Kwan** (Assistant Deputy Minister, Department of Families, Government of Manitoba) in his interview suggests that AI-powered chatbots could be trained with sufficient data and context to generate tailored responses to engage with citizens effectively. Another approach, suggested by other interviewees, involves training AI agents on city data and the residents' information to simulate human behaviour, as an additional supplement for public opinion. This could be useful for decision makers to understand what the public might want in terms of changes to the built environment and public policies. This is similar to how traffic engineers create simulations to represent the movement patterns of a city's population. Using AI agents could merely provide decision makers with another tool to help them make better decisions. However, there is always the risk that these kinds of insights might be interpreted as a replacement for public opinion. As **Santiago Garces** strongly emphasises in his interview, the use of such AI-enabled agents would never replace physical deliberations with real people. These models are only simulations and hence lack the lived experiences of real people, underscoring that generative AI and other forms of technologies cannot substitute public engagement, but merely serve as an additional tool for government workers. to better target and understand local social issues.

Better Public Service Delivery

Generative AI models have the potential to be deployed at the front office of public services to act as an "information assistant" to enable public service delivery to be more accessible and efficient. Typically taking the form of

chatbots, these models could efficiently handle frequently-asked-questions, provide information of public services, and direct queries to relevant call centres, For example, the Ministry of Justice in Portugal has launched the Practical Guide to Justice (GPJ), an AI-powered generative language model which informs citizens and businesses about the tools and services that Justice provides to respond to their needs. Another example can be observed in Singapore, where the government equipped all its websites with an AI-powered information chatbot to handle simple queries as well as municipal reporting. With integration of generative AI into such chatbots, they will be better positioned to understand the analyse the semantics of the queries and requests and subsequently generate personalised responses.

Moreover, generative AI tools have the potential to serve as an 'enabler' for local communities to access information about policies and services affecting them. Instead of navigating various websites to research information about taxes and pension schemes, government policies, or subsidies, residents can easily pose questions on such information from these platforms. In a case-study by Microsoft, generative AI-powered chatbots allow non-English speaking residents in India to access official information reported on government websites and documents through reproducing into their local dialects.

3.2. Roles of Municipalities in Developing Generative AI

Municipalities are grappling with another significant query as they navigate the rapid growth of generative AI: *which role should municipalities play in the development of generative AI tools within their cities? Should this development be solely entrusted to the private sector, or should the municipalities actively participate in developing generative AI models?*

3.2.1. Fostering Collaboration

> Because of municipalities' positions
> as the regulator of cities, they can create
> the environment for innovation within them . . .
> that's not regulating generative AI,
> it is about setting the stage for how
> we want the public and the private sector and
> civil society to interact, it is setting the context.
>
> *Shazade Jameson*
> *Senior Consultant*
> *Tech Governance & Urban AI*

When developing and adopting the generative AI tools, our interviewees suggest that governments and municipalities usually do not possess the inherent capacity (either in terms of technical expertise, manpower or funding), to address all the potential challenges and opportunities surrounding this technology. This requires an ecosystem of actors collaborating to enable the effective development of products and tools. In this way, municipalities have a unique role to play in fostering collaboration and innovation with both the private sector and citizens.

Setting the Stage for Civic Involvement and Participatory Design

Our interviewees suggest that there is potential to utilise generative AI within planning processes to generate different design scenarios based on specific parameters, such as land use zoning or public values. These different options may enable municipalities to facilitate increased public involvement and discussions. Traditionally, the urban design process is led by architects or urban designers who usually present a single output, which may limit the opportunity for dialogue and alternative perspectives. **Anthony Townsend** (Urbanist in Residence, Cornell Tech) highlights the computational capabilities and accessibility of generative AI could enable communities to utilise these tools to propose alternatives or to explore their own ideas. This approach could enhance and empower the role of non-experts and grassroots organisations in the design process, thereby shifting from a planner or municipal-led process to a collaborative and two-way interaction.

By involving the public from the beginning and incorporating their feedback and preferences into design options, governments may prevent scenarios where projects are developed over extended periods without public input or debate, only to be rejected and cancelled later. The adoption of a participatory design builds trust between the government and civic society through demonstrating a commitment to public interests. Moreover, generative AI also holds the potential to enhance our understanding and representation of the built environment. In this regard, **Anthony Townsend** argues that generative AI has potential to generate “stories” - which reflect the public impression and understanding of the benefits and the costs of existing or proposed developments. These outputs could inform both municipalities about the desirability and impact of proposed changes.

However, it is important to note that integrating generative AI alone does not guarantee increased public participation. The responsibility lies with municipalities to foster an environment that encourages and effectively incorporates public opinions into decision-making processes. While leveraging the capabilities of generative AI can certainly facilitate and optimise this collaboration, it is essential for municipalities to take proactive steps in setting grounds and implementing strategies that foster meaningful public participation.

Fostering Public and Private Sector Collaboration

The public sector usually lacks the capacity and internal knowledge and skills required to develop generative AI tools and models independently. Thus, the public sector needs to take advantage of the knowledge and capabilities of the private sector. The role of individuals who could bridge between municipalities and the private sector becomes particularly important, enabling municipalities to decipher what aspects of generative AI may be useful for them. **Maria Mamoura** (Director, Bryden Wood) asserts that the development of AI tools will predominantly take place in the private sector, rather than governments building capacity in-house, as governments can be financially constrained in investing in the development of internal teams to explore these technologies. The fostering of public-private partnerships would thus be crucial in aiding private companies to provide the necessary hardware and model training capabilities, while the government contributes through regulation and the plethora of data that it has access to.

3.2.2. Responsible Use of Technology adopted from Private Sector

> The process has to be circular.
> It has to involve humans in the loop all the time,
> even with the price of slowing the process down
> a bit and getting less insights per second.
>
> *Ariel Noyman*
> *Urban Scientist, MIT Media Lab*

While governments and municipalities are gradually incorporating AI tools into their workflows and data analysis processes, the integration of generative AI is still in its nascent stages. Hence, it is crucial for the city municipalities to recognize that they have a proactive role in the development and implementation of generative AI technology, rather than simply relying on the private sector to develop tools and purchasing for use later. **Emily Binet Royall** (Smart Cities Administrator, City of San Antonio) emphasises the potential pitfalls that can arise when municipalities blindly adopt and apply private AI tools to public sector use cases without closely monitoring their development and understanding their functionality. Many companies build the tools with a global scale in mind, rather than catering specifically to the distinct characteristics, cultural expectations, and behavioural patterns of individual cities. Additionally, developers of AI tools cannot always guarantee the explainability of their data models, as explainability is a well known trade-off with reliability and accuracy of AI generated outcomes. Data models are also often proprietary, creating a

challenge for companies to both protect their IP and facilitate transparency. Ultimately, one-size-fits-all approaches may prove ineffective in diverse urban contexts, hence imposing risk of the loss of individual agency by municipalities.

Therefore, municipalities should actively participate in the training process by testing the tool's responses and providing their own data to ensure it is applicable for local use. **Paula Boet Serrano** (Project Manager, Barcelona City Council) also stresses the necessity of paying attention to the limitations of the models, such as their potential to generate misleading or false information and hallucinations (refer to Section 3.4.1), and to establish safeguards accordingly. She suggests incorporating human evaluation and feedback into the training process at every stage. This human involvement ensures transparency, accountability, and a critical understanding of the tools and technologies being implemented. Furthermore, by involving humans in evaluating the results and providing feedback, municipalities can iteratively improve AI models. Moreover, **Shazade Jameson** (Senior Consultant Tech Governance & Urban AI) highlights that it is equally important for municipalities to create test beds to allow for comprehensive testing and assessment, ensuring their reliability and effectiveness in practical scenarios. This approach can help in ensuring transparency, accountability, and certification of AI tools for public sector use. Lastly, municipalities may influence the development of Generative AI through local regulations and governance frameworks, which will be explored further in the subsequent section.

3.2.3. Development of AI Governance Frameworks Prior to Implementation

> Before we get to a point where we could rely on generative AI to be a part of that decision making process, we will need to develop an AI governance that considers the ethical as well as the liabilities and legal implications
>
> .
>
> *Ernest Kwan*
> *Assistant Deputy Minister, Department of Families, Government of Manitoba*

Our interviewees strongly emphasise the importance of establishing robust governance frameworks and standards for generative AI before incorporating it into decision-making. These frameworks should address multiple ethical and legal considerations (refer to Section 3.4.2) and provide guidance to ensure the responsible implementation of generative AI models, with the

aim of minimising potential risks to the public. The process of establishing governance includes, but is not limited to, identifying who is responsible for the development, deployment, and maintenance of these systems, as well as establishing procedures and protocols for monitoring and evaluating their performance. An important consideration here is that such governance protocols should be applicable and relevant to the specific local context of the city. According to the concept of AI Localism by Verhulst et al. (2021), it was strongly emphasised that national and federal legislations often lack critical consideration of local context, and the role of municipalities emerged to fill such gaps. Calibrating AI policies for local conditions and managing them within tightly defined geographic regions would enable policymakers to better understand the stakeholders and tradeoffs involved, hence better likelihood of creating positive feedback loops, resulting in increased effectiveness of the policies. This is why municipalities need to better assess their local value with their citizens to regulate those technologies. **Shazade Jameson** highlights through her interview that regulatory frameworks should be **principle-driven** - they should not be tied to specific technologies, but rather the emphasis should instead be on the impact that these technologies are expected to achieve. For example, the City of Boston has taken a pioneering step by publishing an interim guideline for the use of generative AI tools. This development comes in the wake of enabling Google Bard enterprise access for 10,000 public servants within the city. Moreover, It is also crucial for these frameworks to align with existing regulations and policies concerning AI and other technologies, such as the EU AI Act, especially if the municipality falls within the EU.

3.3. Skills and Workflows

Developing generative AI tools and integrating them into municipal workflow is neither a straightforward and literal process, due to both the intricate nature and the nascent stage of this technology, which could vary across individual cities. Concurrently, It also necessitates new skills to be developed and employed to navigate this emerging field. This section discusses skills and the development of procedures within urban governance that have potential to facilitate the successful and efficient development and integration of generative AI in municipal settings.

3.3.1. Internal Capacities to be Built in Municipal

Problem Definition and Engineering

A city is composed of many layered and interconnected components, including urban infrastructure such as roads and buildings, scales of governance such as council, district and municipal governments, and individuals and citizen organisations. Due to this complexity, many urban issues are often broadly-defined and the government often lacks the ability to identify and articulate a clear logic of change for these issues. The development of generative AI

for application within cities would thus require clear problem definition and engineering to break down these issues and define them.

Municipalities may begin by problem formulation by which an urban problem is clearly defined and its causal implications are explored as well as the potential application of generative AI to address it. Problem definition should be followed by problem engineering - the process of breaking down complex issues into manageable components. This is particularly important when tackling urban problems which are often multifaceted and complex. This process could include a wide array of action steps, including but not limited to, analysing factors that contribute to both the creation of the problem and the proposed solution, and determining and continuously evaluating the most appropriate strategies for addressing them.

Critical Understanding of Technology and Tools

Government needs to have those individuals who understand how the tools and technology works, even if they do not directly create them. **Maria Mamoura** suggests that the municipalities may not necessarily need to hire in-house developers to develop these tools, however, it is crucial to have personnel (specialists, consultants etc.) with technical expertise who can critically understand and assess how the tools and technology works. Municipalities will need to build this in-house capacity as these specialists will not only be required to understand the tools, but also serve as conduits for disseminating this knowledge to other members of the municipal workforce.

Ethical Regulations and Fair Contracts

Paula Boet Serrano argues that the municipality should establish ethical regulations, protocols and guidance in development and implementation of generative AI-powered public services. Furthermore the development of contracts that govern these partnerships is imperative. These contracts should be carefully crafted to encompass fair clauses and transparency requirements, as well as technical and quality assurance to ensure that the tools developed under such contracts are of good quality and contribute to the public interest. By ensuring that the contractual agreements are fair and binding, municipalities are more likely to uphold accountability, transparency, and the overall welfare of the communities they govern in the implementation of AI technologies in public services. **Paula Boet Serrano** also mentions that the City of Barcelona is collaborating with the European Commission and other cities to develop these standard algorithmic clauses.

3.3.2. A Step by Step Process

The integration of generative AI into critical decision-making processes for municipalities will likely be a gradual and extended process rather than quick penetration and rapid adoption. First and foremost, the high-risk nature and black box characteristics of AI algorithms necessitate multiple rounds of testing

to ensure the usefulness and applicability of the model's outputs for specific problem domains. Whether developed internally within the government or through public-private partnerships, any AI-powered models should undergo rigorous internal testing against specific use cases or testbeds. Additionally, **Shazade Jameson** highlighted that municipalities must also consider other factors, such as digital maturity of the city's cyber-physical infrastructure, and build the necessary capacity, if required, during the implementation of these models. As a result, the overall process becomes lengthy and gradual, accounting for the complexities and considerations involved in successfully integrating generative AI into municipal decision-making processes.

> Will it quickly penetrate the core of how cities are planned, analysed and governed - I think it's going to be a long time
>
> *Ernest Kwan,*
> *Assistant Deputy Minister, Department of Families,*
> *Government of Manitoba*

Ernest Kwan recommends, in his interview, a phased approach from a closed to an open-system, when implementing a generative model. This gradual and phased implementation allows for collaboration opportunities and knowledge sharing before making the model available to the general public. When implementing their own generative AI service, municipalities should not only mine their own data, they also need to ingest other sources of information - either data or technical capacity from private sectors as well as practices and protocols from other municipalities around the world. By leveraging these diverse sources, municipalities can enhance the effectiveness of their models through open collaboration.

There is also a need to develop organisational capacity for accountability, including establishing feedback mechanisms that enable continuous monitoring of AI implementation and usage within the organisation over time. **Ernest Kwan** also emphasises the significance of starting with pilot projects. By conducting multiple pilot testing phases, organisations can gradually refine and fine-tune the generative AI models. This incremental approach offers the municipalities the opportunity to learn from initial implementations, address errors, and build feedback loops and spaces for discussion along the way. The success demonstrated in pilot projects serves as a foundation for broader adoption, ensuring that the implementation of generative AI is based on proven effectiveness and suitability for the intended purposes.

3.3.3. Containing Risks and Sandboxing

Containing risks within specific boundaries and limiting exposure to other stakeholders is important. This approach allows experimentation and innovation to take place within a controlled environment, while allowing the municipal and private sector stakeholders to maintain a balance between technological advancement and risk mitigation. As an example, **Shazade Jameson** uses the concept of "sandboxing". A sandbox is essentially a controlled space within the market where a subset of organisations, startups, or individuals can test their technology with relaxed regulations. By providing this controlled environment, private firms can explore and validate innovative ideas and technologies with multiple trials and limited exposure of risk to the public. However, the success of sandboxing hinges on the negotiation of expectations for entry and exit of the sandbox, which relies on the social relationships and collaboration that was highlighted earlier.

3.4. Challenges and the Future of Generative AI for City Governance

> We are also exploring how this kind of technology
> can be used in the right way
> because we are a public administration.
> So we need explainability and robustness,
> but especially explainability
>
> *Paula Boet Serrano*
> *Project Manager, Barcelona City Council*

The utilisation of generative AI in municipal governance and public services requires careful consideration of potential challenges that accompany this technology due to its high-risk nature and potential impact on the general public. **Anthony Townsend** suggests that the application of generative AI in domains such as municipal decision-making, which involves crucial judgments and significant accountability, may give rise to obstacles that require the development of robust governance frameworks. In light of ideas such as this, this section reflects on challenges around the use of generative AI and potential mechanisms to foster safer outcomes.

3.4.1. Technical Errors

Hallucination

Generative AI hallucinations occur when the model generates responses that are fabricated partly or entirely in a way that may look correct and plausible, but is not actually valid and factual. For instance, if a user asks ChatGPT to find a news report on a specific issue, the model may generate a response that includes a fully-fledged news article, web links and even citations to external publications. However, upon verification of the sources, it may be the case that those articles are non-existent and are merely fabricated content. These hallucinations can be problematic, particularly if these models were to be employed in public engagement and decision-making.
AI companies like OpenAI have already explored strategies to mitigate such hallucinations. OpenAI recently launched a research paper that outlines their new "process supervision" approach (Lightman et al., 2023) in training models to combat hallucinations. In a similar vein, the municipalities should also carefully incorporate fact-checking mechanisms in their own development process.

Technical Debt

> "As the use of generative AI expands, we're going to accumulate a lot of technical debt and I think that's a really bad thing for governments to casually wander into."
>
> *Anthony Townsend*
> *Urbanist in Residence, Cornell Tech*

Technical debt refers to the risk accumulated throughout the development cycle, which later needs to be refactored when shortcuts or suboptimal solutions are taken during the development process. It is often the result of prioritising speedy delivery over perfecting the algorithms to minimise errors. **Anthony Townsend** advocates that the municipalities who may develop generative AI applications should evaluate their models against the level of technical debt. This way, municipalities can better assess the long-term sustainability and quality of their models.

Reproducibility

One common challenge in generative AI is that it is relatively easy to generate

content, but reproducing those results reliably is a difficult task. Reproducibility refers to the ability of repeatedly running the same prompts on the algorithm and obtaining the same (or similar) results on each run. **Anthony Townsend** emphasises the issue of reproducibility and the ability to generate consistent outputs as an ongoing concern within the field. He points out that many of the current processes heavily rely on trial and error methodologies, hence lacking a systematic understanding of how end results and outputs are produced. However, this issue is still universal, with no solution readily available to completely eliminate it. This necessitates rooms for further investigation and exploration by the tech researchers and experts, meanwhile city municipalities should actively monitor any developments and contribute if applicable.

3.4.2. Ethical Considerations

Data quality and biases

Many AI tools developed externally by private firms usually are trained on public datasets (for examples, DALL·E 2 is trained on hundreds of millions of captioned images publicly available on the internet.) This raises several issues in terms of data quality and potential bias. Most public datasets lack transparency around dataset design and collection procedures as these information are often not publicly disclosed. Moreover, many public datasets are either labelled (e.g., “apple” for all pictures of apples) or require manual labelling by model developers before training the model. In both cases, presence of biassed labels or ethically questionable annotations can affect the quality of datasets as well. Lastly, imbalance of representation and lack of variety. A research study by Yang et al., (2019) claimed that the best algorithm wouldn’t work well if the data it learned from didn’t reflect the real world.

Transparency

Algorithmic transparency registers and portals, such as the ones established in Amsterdam and Helsinki, play a significant role in building transparency of the algorithms utilised by the government. These portals disclose and outline critical information of the algorithms in use, such as how the model is trained, how its predictions and outputs are used, and how potential biases or risks are addressed. Moreover, these registries offer citizens the opportunity to provide feedback on the algorithms employed by their local government. Such initiatives have the potential to empower citizens, enabling them to evaluate, examine, and scrutinise how the government integrates generative AI tools.

Interpretability

Despite that many generative AI models on market today are user-friendly, they still lack interpretability, which is the user’s capacity to comprehend how the model operates and how it generates output. This lack of interpretability

raises significant concerns. As a result, the municipality may not be equipped with a complete understanding of a decision that was generated from an AI model. **Paula Boet Serrano** brought out the need for the municipality to **explain** the decision output of the model to the public so that they can adopt it. However, the inherent opaqueness of generative AI limits the explainability of the decision-making process and, consequently, inhibits the ability to unpack the model algorithms for the public at large. Ultimately, the fundamental lack of interpretability underscores the need for cautious evaluation and the establishment of **mechanisms** to ensure transparency and responsible implementation. This indicates that generative AI may not be used independently as the sole decision-maker or to carry out the decision-making process on its own.

Building Civic Trust in Technology

Building civic trust in the generative models that municipalities intend to deploy is of critical importance. **Ernest Kwan** highlights that a single inappropriate or mishandled response can have significant consequences, impacting the government's reputation and credibility as well as the well-being of the residents. Without proper governance, the installation of new technology can erode public trust and confidence in government institutions (Gordon and Mugar, 2020). Hence, the implementation of generative AI for government decision-making should aim to incorporate robust check-and-balance mechanisms and comprehensive legislative frameworks throughout the entire lifecycle to ensure its relevance, applicability, and value for the city and its residents, whilst mitigating potential exposure of risk.

Privacy and Confidentiality

Generative models, particularly language models, have capacity and ability to memorise data (Biderman et. al, 2023) from the training datasets. This raises a significant concern as many language models store user input data for training purposes, unless the user explicitly opt-out, to continuously expand its training dataset and improve the performance. As a result, personal information or other confidential information inputted in the prompt may be leaked to the public at large or extracted by malicious actors. Such incidents would not serve to build civic trust of the public since their privacy and personal information might be at risk of disclosure. A similar case was reported recently wherein Samsung Electronics Co. banned the use of generative AI platforms following the incident of an employee accidentally leaked internal source code by uploading it to ChatGPT. **Santiago Garces** also emphasises this concern in his generative AI guide for Boston City, advising against including confidential information in the prompt and emphasising the need to ensure the deletion of such information from notes or other inputs.

3.4.3 Incorporating Expert Knowledge, Human Oversight and Setting Benchmarks

Paula Boet Serrano underscores the need for governments to communicate openly with citizens to identify their concerns and expectations on new technologies like generative AI, as well as, to unpack and explain how these AI-powered tools can impact the residents. **Ernest Kwan** suggests that a possible approach is to create and facilitate a democratic mechanism for evaluating the reliability of generative AI outputs, such as consensus-based or vote-based systems where the local residents (who are primary users) can rate the reliability of the data or responses as an internal-trial. Moreover, **Ernest Kwan** also suggests that academics, research agencies, and other domain experts can become "bridges" in channelling public sentiment and insights into feedback mechanisms inside the municipality. Through seeking guidance from these experts, the municipal authorities can leverage their insights to either improve the tools or drive adoption of the tools.

3.4.4. Generative AI as a Political Tool

> Those who are developing and evaluating these tools for public sector use must think critically about how political these tools can be.
>
> *Emily Binet Royall*
> *Smart Cities Administrator, City of San Antonio*

Municipalities who are intending to explore generative AI for the public sector must critically analyse and account for the potential political and social risks and complications. **Emily Binet Royall** asserts that municipalities need to recognise that generative AI tools can be political in nature. Cities may often adopt technological solutions that were evaluated only through a technology lens, failing to acknowledge their other social and political aspects.

If these models contain biases or discriminatory patterns in their outputs, it could potentially result in unfair decisions that exacerbate existing inequalities. **Emily Binet Royall** provides a real-life example of a municipal department purchasing a solution to automate the analysis of street conditions. However, the results of this automated analysis indicated that more funding should be allocated to the wealthiest districts rather than the poorer ones. And the officers responsible for using the tool were unable to understand or explain how this output was generated. This scenario highlights how a single incorrect decision stemming from an AI model could significantly impact the lives of many, as resources or policy actions may not be directed to those residents who are most in need, as a result.

Anthony Townsend also echoed a similar sentiment that this technology could be weaponised by powerful interests to shape the public discourse in their favour, hence undermining the public interest. Generative AI models can also be used for malicious purposes, such as creating disinformation, driving misinformation campaigns, and impersonating individuals as **Zhongwen Huang** (Director - Smart City Projects Office, Smart Nation and Digital Government Office Singapore) highlighted during his interview. Machine-generated images and texts can invoke shocks among the public or even sway the public opinions in desired ways. Such fake news and misinformation are not new issues brought about by generative AI. However, the computational capabilities of generative AI enables anyone to mass-produce such fabricated contents at minimal time and costs. This can be potentially problematic from the security and governance perspective. According to Sarah Kreps, a Cornell professor, guardrails around generative AI are emerging but the real challenge lies in the new and dynamic nature of the technology that its projection is almost unpredictable.

An opportunity that **Emily Binet Royall** highlighted is that the city can prepare its future workforce to safeguard the potential political risks of these technologies. She mentioned that municipalities should facilitate internships within the city government for young people who have undergone engineering courses, data science programs, or machine learning training. By exposing the future tech workforce to how the tools they develop are applied as political tools for public good, they can gain firsthand experience and insights into the impact and consequences of their work.

5. RECOMMENDATIONS

The transformative potential of generative AI in managing and planning cities is evident. It has the capability to aid in understanding and visualising the impact of policy changes and development plans, serving as a valuable foundation for productive dialogue and debate between municipal authorities, the public and private sector. Simultaneously, the application of generative AI in the public sector brings a number of risks and pitfalls, as specifically outlined in Section 3.4 of the Findings. This is one of the greatest challenges of this technology, striking a careful balance between risk mitigation and harnessing its capabilities for public good. To assist with this task, we have outlined a set of recommendations in the form of high-level action steps for reference by municipal governments.

(1) Generative AI as an Enhancer of Existing Processes

It is unlikely that generative AI will replace current processes within urban governance, but it could serve as a valuable enhancer for them. For example, our research suggests that generative AI tools will not replace government workers, but it does have the potential to enhance their productivity. It will not replace architects and urban planners, but it could provide them with more tools to visualise and explore scenarios as well as better plan for future conditions based on historical data. In line with this thinking, neither is a generative AI model a substitute for true public participation and community engagement, but it does have the potential to enhance it through both empowering government officers to better understand the diverse public and allowing citizens new ways of engaging with urban developments and policy.

(2) Principle-Driven Policy and Iterative Regulatory Mechanisms

The setting of clear regulations governing the development and utilisation of generative AI for public sector use cases is imperative. Whilst municipalities may not directly develop these technologies and tools in-house they should maintain an active role throughout the entire process. It is also important to avoid the adoption of ready-made, “one-size-fits-all” solutions from the private sector, ensuring that the unique needs and context of each municipality are taken into account. Municipalities should also actively monitor the recommendations made by standards-setting bodies (federal, national, and regional) to continuously improve and modify the regulations surrounding the use of generative AI in public sector settings.

(3) Establishing Robust Channels for Collaboration and Innovation

Municipalities should establish robust channels and opportunities for collaboration and innovation with both civil society and the private sector. It is crucial to identify and involve relevant stakeholders such as public officials,

civil servants, urban operators, technology vendors, and city residents, considering their diverse interests. This inclusive approach ensures meaningful and appropriate engagement of each party throughout the process.

(4) Civic Good over Economic Growth

Private firms, startups, and developers involved in public-private partnerships for generative AI development must recognise that the products and tools they create may not solely focus on commercial scalability. Instead, they should prioritise civic good over economic growth.

(5) Building Internal Capacities and Frameworks for accountability and responsibility

There is a need for municipalities to develop organisational capacity for accountability, including establishment of feedback mechanisms to monitor AI implementation and usage within the organisation over time. This could be facilitated through the development of frameworks to clearly identify the sharing of responsibility, risks and accountability among different stakeholders (municipalities, private sector partners, developers, etc.) throughout the entirety of the development lifecycle, from conception, to deployment and maintenance. Municipalities would also need to set up appropriate guardrails and boundaries from the beginning to ensure that the model or tool developed does not deviate from the intended application or produce any malicious or destructive outcomes. The public sector is high risk as it has a direct impact on all citizen's lives, therefore municipalities need to operate within robust frameworks that ensure accountability and responsibility for the wellbeing of all citizens.

(6) Institutional Challenges

For municipalities to effectively adopt generative AI, there are two main institutional challenges. Firstly, the municipalities should carefully consider, plan and prepare the necessary processes, skills and workflows to allow for successful implementation of this technology. Secondly, municipalities need to shift from old ways of working to new ones that better fit with the new workflows enabled by technology integration and innovation. This might disrupt the existing internal hierarchies of the institution by creating new roles or redefining current ones through digital transformation. It is important for the municipalities to embrace these changes and provide support, training and resources to the workforce to adapt to these new roles and responsibilities. Municipalities need to be open for innovation, development and evolution at an institutional level. Changing the culture and mindset of the institution to embrace innovation and adapt to new practices is a continuous effort, requiring long-term commitment.

REFERENCES

Alshuwaikhat, H. M., & Abussaud, L. Z. (2023). Generative Urban Design for Smart, Sustainable Resilience Cities. 42nd Annual Conference of the International Association for Impact Assessment.

Anderson., J. & Rainie., L. (2018). Artificial Intelligence and the Future of Humans. Pew Research Center: Internet, Science & Tech. https://www.pewresearch.org/internet/2018/12/10/artificial-intelligence-and-the-future-of-humans/

Baxter, K., & Schlesinger, Y. (2023, June 6). Managing the risks of Generative AI. Harvard Business Review. https://hbr.org/2023/06/managing-the-risks-of-generative-ai

Berns, S., & Colton, S. (2020). Bridging Generative Deep Learning and Computational Creativity. Proceedings of the 11th International Conference on Computational Creativity (ICCC'20), 406–409.

Bertuzzi, L. (2023). Ai act moves ahead in EU Parliament with Key Committee vote. Euractiv. https://www.euractiv.com/section/artificial-intelligence/news/ai-act-moves-ahead-in-eu-parliament-with-key-committee-vote/

Biderman, S., Prashanth, U. S., Sutawika, L., Schoelkopf, H., Anthony, Q., Purohit, S., & Raf, E. (2023). Emergent and predictable memorization in large language models. arXiv preprint arXiv:2304.11158.

Bird, E., Fox-Skelly, J., Jenner, N., Larbey, R., Weitkamp, E., & Winfield, A. (2020). The ethics of Artificial Intelligence: Issues and initiatives. European Parliamentary . https://www.europarl.europa.eu/RegData/etudes/STUD/2020/634452/EPRS_STU(2020)634452(ANN1)_EN.pdf

Cai, Q., Ma, M., Wang, C., & Li, H. (2023). Image neural style transfer: A Review. Computers and Electrical Engineering, 108, 108723. https://doi.org/10.1016/j.compeleceng.2023.108723

Chui, M., Hall, B., Mayhew, H., Singla, A., & Sukharevsky, A. (2022). The state of AI in 2022-and a half decade in Review. McKinsey & Company. Retrieved from https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review

Chui, M., Roberts, R., & Yee, L. (2022). Generative AI is here: How tools like ChatGPT could change your business. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/generative-ai-is-here-how-tools-like-chatgpt-could-change-your-business

Clarke, L. (2022). Is generative AI really a threat to creative professionals?. The Guardian. https://www.theguardian.com/technology/2022/nov/12/when-ai-can-make-art-what-does-it-mean-for-creativity-dall-e-midjourney

REFERENCES

Cremer., D. D. & Kasparov G., (2021). AI should augment human intelligence, not replace it. Harvard Business Review. https://hbr.org/2021/03/ai-should-augment-human-intelligence-not-replace-it

Esling, P., & Devis, N. (2020). Creativity in the era of Artificial Intelligence. https://doi.org/10.48550/arXiv.2008.05959

Feldstein, S. (2019). The Global Expansion of AI Surveillance. Carnegie Endowment for International Peace. https://carnegieendowment.org/files/WP-Feldstein-AISurveillance_final1.pdf

Fitzgerald, M. (2016). Data-Driven City Management. MIT Sloan Management Review.

Gioti, A. M. (2021). Artificial Intelligence for Music Composition. Handbook of Artificial Intelligence for Music, 53–73. https://doi.org/10.1007/978-3-030-72116-9_3

Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. (2014) Generative Adversarial Nets. Advances in Neural Information Processing Systems, 27, 2672-2680.

Grace, K., Salvatier, J., Dafoe, A., Zhang, B., & Evans, O. (2018). When will AI exceed human performance? evidence from AI experts. Journal of Artificial Intelligence Research (AI and Society Track). https://arxiv.org/abs/1705.08807

Hämäläinen, M. (2020). A Framework for a Smart City Design: Digital Transformation in the Helsinki Smart City. In: Ratten, V. (eds) Entrepreneurship and the Community. Contributions to Management Science. Springer, Cham. https://doi.org/10.1007/978-3-030-23604-5_5

Hinkley, Sara. (2022). Technology in the Public Sector and the Future of Government Work. Berkeley:UC Berkeley Labor Center. https://laborcenter.berkeley.edu/technology-in-the-public-sectorand-the-future-of-government-work/

Huzaifah, M., & Wyse, L. (2021). Deep generative models for musical audio synthesis. Handbook of Artificial Intelligence for Music, 639–678. https://doi.org/10.1007/978-3-030-72116-9_22

Jaupi, J. (2022). Could artificial intelligence really wipe out humanity? New York Post. https://nypost.com/2022/02/20/could-artificial-intelligence-really-wipe-out-humanity/

REFERENCES

Jin, L., Tan, F., & Jiang, S. (2020). Generative adversarial network technologies and applications in computer vision. Computational Intelligence and Neuroscience, 2020.

Kandt, J., Batty, M., (2021). Smart cities, big data and urban policy: Towards urban analytics for the long run. Cities 109, 102992. https://doi.org/10.1016/j.cities.2020.102992

Kitchin, R., Lauriault, T., (2014). Small Data, Data Infrastructures and Big Data. https://doi.org/10.2139/ssrn.2376148

Li, Z., & Wu, Y. (2021). The effectiveness of image augmentation in breast cancer type classification using Deep Learning. 2021 3rd International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI). https://doi.org/10.1109/mlbdbi54094.2021.00134

Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I., & Cobbe, K. (2023). Let's Verify Step by Step. arXiv Preprint arXiv:2305.20050. https://doi.org/https://doi.org/10.48550/arXiv.2305.20050

Lorica, B., & Loukides, M. (2016). What is Artificial Intelligence? O'Reilly Media. Retrieved May 3, 2023, from https://www.oreilly.com/radar/what-is-artificial-intelligence/

Marshan, A. (2021) Artificial intelligence: Explainability, ethical issues and bias. Ann Robot Automation 5(1): 034-037. DOI: 10.17352/ara.000011

McKendrick, J. (2023). Artificial Intelligence opens up the world of Financial Services. Forbes. https://www.forbes.com/sites/joemckendrick/2023/02/24/artificial-intelligence-opens-up-the-world-of-financial-services/?sh=6c472938b0ee

Nizzad, A. R. M., Sameer, U. M., Suhath, S. M., Nafrees, A. C., Rankothge, W. H., Kehelella, P. H., & Mansoor, C. M. M. (2021). Internet of things based automatic system for the traffic violation. 5th International Conference on Electrical, Electronics, Communication, Computer Technologies and Optimization Techniques (ICEECCOT). https://doi.org/10.1109/iceeccot52851.2021.9708060

O'Sullivan., M. (2021). Will AI Take Over The World? Forbes. https://www.forbes.com/sites/mikeosullivan/2023/03/11/will-ai-take-over-the-world/?sh=608589be11df

REFERENCES

Paetzhold, M. (2022). Will dall-e the AI artist take my job?. Intelligencer. https://nymag.com/intelligencer/article/will-dall-e-ai-artist-take-my-job.html

Pfizer. (2018). Artificial Intelligence: On a mission to Make Clinical Drug Development Faster and Smarter. Pfizer. https://www.pfizer.com/news/articles/artificial_intelligence_on_a_mission_to_make_clinical_drug_development_faster_and_smarter

Purohit, A. (2023). Ai, ML, DL, and generative Ai face off: A comparative analysis. Synoptek. https://synoptek.com/insights/it-blogs/data-insights/ai-ml-dl-and-generative-ai-face-off-a-comparative-analysis/

Radford, A., Narasimhan, K., Salimans, T. & Sutskever, I. (2018) Improving Language Understanding by Generative Pre-Training.

Russell, S. J., & Norvig, P. (2010). In Artificial Intelligence: A modern approach. preface, Prentice-Hall.

Stanford University. (2022). Artificial intelligence (AI) adoption worldwide 2021, by industry and function [Graph]. In Statista. Retrieved from https://www.statista.com/statistics/1112982/ai-adoption-worldwide-industry-function/

Toews, R. (2022). 4 predictions about the Wild New World of text-to-image ai. Forbes. https://www.forbes.com/sites/robtoews/2022/09/11/4-hot-takes-about-the-wild-new-world-of-generative-ai/?sh=5d82ff0e13d9

Torkington, S. (2023). The golden age of ai: Why chatgpt is just the start. World Economic Forum. https://www.weforum.org/agenda/2023/03/generative-ai-chatgpt-machinelearning/

United Nations Population Division. (n.d.). Urbanization. United Nations. https://www.un.org/development/desa/pd/content/urbanization-0

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need. Advances in Neural Information Processing Systems 30 (NIPS 2017). https://doi.org/https://doi.org/10.48550/arXiv.1706.03762

Verhulst, S., Young, A., & Sloane, M. (2021). The AI Localism Canvas. Informationen zur Raumentwicklung, 48(3), 86-89. https://ssrn.com/abstract=3946503

REFERENCES

Wang, D., Lu, C.-T., & Fu, Y. (2022). Towards Automated Urban Planning: When Generative and ChatGPT-like AI Meets Urban Planning. ACM Trans. Spatial Algorithms Syst., Vol. 1, No. 1, Article 1. https://doi.org/https://doi.org/10.48550/arXiv.2304.03892

Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Tao, A., Kautz, J., & Catanzaro, B. (2018). High-resolution image synthesis and semantic manipulation with conditional gans. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr.2018.00917

World Bank. (2023). Urban Development Overview. https://www.worldbank.org/en/topic/urbandevelopment/overview

Wu, A. N., Stouffs, R., Biljecki F. (2022). Generative Adversarial Networks in the Built Environment: A Comprehensive Review of the Application of GANs across Data Types and Scales.Building and Environment, 223: 109477

Yang, K., Qinami, K., Fei-Fei, L., Deng, J., & Russakovsky, O. (2020). Towards fairer datasets. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency.* https://doi.org/10.1145/3351095.3375709

URBAN AI

Urban AI is a Paris-based think tank dedicated to the emerging field of "Urban Artificial Intelligence." Urban AI seeks to generate a holistic body of knowledge on urban artificial intelligence by federating and collaborating with a growing, global community of researchers, public servants, start-ups, and urban subject matter experts, who work at the intersection of cities and technology. Together, we carry out multidisciplinary projects to better understand and assess the impacts of artificial intelligence on urban life and vice versa.