

**Early Exploration of GPT-4's Abilities
in the Context of Urban Planning and Design in Singapore**

Tuan-Yee CHING and Alvin CHUA

Urban Planning and Design Technology Centre of Excellence (URBEX)
Design & Planning Lab, Urban Redevelopment Authority of Singapore

Author Note

Tuan-Yee Ching and Dr. Alvin Chua are from the Design & Planning Lab, Urban Redevelopment Authority (URA) Singapore.

URBEX (Urban Planning & Design Technology Centre of Excellence) aims to build up strategic and critical Science, Technology and Engineering capabilities to advance urban planning and design for Singapore. We drive innovation, encourage experimentation and sharing of lessons learnt, in collaboration with partner organisations.

The URBEX working paper series serves as a platform to share insights from on-going studies with the wider urban planning community. The overall intent is to rally communities of interest around strategic and emerging subject areas of relevance to urban planning & design technology.

Correspondence concerning this article should be addressed to: urbex@ura.gov.sg

Abstract

This paper shares observations from an early exploration of the abilities of OpenAI's latest large language model (LLM), GPT-4, in the context of urban planning and design in Singapore. GPT-4 excels in articulating urban planning and design concepts and strategies, with structured analyses and arguments resembling those of an experienced urban planner. It is able to suggest technical analysis methodologies, including steps and code for running Geographic Information System (GIS) software and generating simple 2D and 3D diagrams. However, GPT-4 struggles with providing links to publications and lacks the ability to extract and present geospatial data. It also falls short in responding to Singapore-specific domain knowledge and hallucinates on some facts. This points to the need to train it with domain-specific information. GPT-4 can already assist urban planners in articulating well-rounded analyses and reaping productivity gains on daily tasks. However, it does not replace the need for urban planners to continue building competencies, knowledge, and understanding of cross-domain issues, data, tools, and analytics. Planners themselves need to fact-check, reason critically and apply knowledge while working alongside GPT-4.

Keywords: large language models, GPT-4, urban planning, urban design, Singapore, planning competency

Early Exploration of GPT-4's Abilities in the Context of Urban Planning and Design in Singapore

Highlights:

- GPT-4 excels in articulating urban planning and design concepts and strategies. Its structured arguments resemble those of a seasoned urban planner. It can already assist planners in articulating well-rounded analyses and reaping productivity on daily tasks. It suggests technical analysis methodologies, including steps for running GIS software and generating simple 2D and 3D diagrams.
- GPT-4 struggles with providing functional internet links and lacks the ability to extract geospatial data. It falls short in responding to Singapore-specific domain knowledge and hallucinates on facts. This points to the need to supplement it with domain-specific information.
- It does not replace the need for urban planners to continue building competencies, knowledge, and understanding of cross-domain issues, data, tools, and analytics. It is still up to human planners to fact-check, reason, critically think and apply knowledge.

Introduction

GPT-4, OpenAI's latest large language model (LLM) was released in Mar 2023, just four months after the launch of ChatGPT (3.5) which took the world by storm. This paper shares findings from an early exploration of GPT-4's abilities in the context of urban planning and design in Singapore.

Amidst the growing interest in LLMs, the American Planning Association (APA) recently published a Quick Note on "ChatGPT: Implications for Planning" (Daniel, 2023). The publication outlined the potential for ChatGPT to be the "next level of office AI assistant" in drafting emails and reports, performing simple coding tasks, and creating workflows in software applications such as ArcGIS and CAD. However, the document also warned that "ChatGPT has no sense of 'truth'" and

recognised the need to provide more machine-readable information to close the gap in Artificial Intelligence (AI) training data. On the one hand, AI-generated content can lower barriers to expressing ideas and views on planning matters, yet on the other hand, can be a source of disinformation.

OpenAI did not specify the actual training data used to train GPT-4, although this is widely assumed to be publicly-available Internet data. Based on OpenAI's report, GPT-4 outperforms ChatGPT in several exams and benchmarks (OpenAI, 2023). Bubeck et al (2023), in their investigation of GPT-4's abilities, also concluded that compared to previous AI models, GPT-4 exhibited more general intelligence. With the rapid development of LLMs and the growing number of software applications incorporating such models, there is urgency to assess GPT-4's abilities and readiness for use in specialist domains like urban planning.

Background and Related Work

AI is a field of computing associated with developing systems capable of performing tasks that require human intelligence. The use of AI in urban planning can be traced back to the 1960s and 70s, where early attempts were made to use mathematical models for land use-transport and urban growth simulations (Batty, 1976). These models provided urban planners with insights into the dynamics of urban systems and gave them a means to evaluate the impact of new plans and policies. However, these models remained largely theoretical due to limited data and computational resources.

With the emergence of more powerful computers in the 1980s and 90s, attempts were made to develop expert systems for planning decision support (Ayeni, 2003). Such systems were driven by rule-based logic models and were created to

accelerate workflows through automation. These systems made specialist knowledge available to a wider user base but were costly to create as significant amounts of time were required to model the logic into applications.

As data collection methods and technology improved in the early 2000s, there was a shift towards more data-driven approaches for decision-making. Machine learning algorithms and big data analytics were incorporated with more traditional analytical methodologies to study and model the impact of trends on urban development.

Deep learning (LeCun, 2015) has become increasingly prominent in recent years due to its ability to train highly performant models on large datasets. Generative Pre-Trained Transformer (GPT) is a type of LLM that uses deep learning techniques to process and generate natural language. GPT-4, the largest LLM released by OpenAI was reported to exhibit remarkable capabilities in a variety of tasks, including abstraction, comprehension, vision, coding, mathematics and understanding of human motives amongst others (Bubeck, 2023). Thus far, the use of AI technologies in urban planning has been limited to low-level operations like computation, pattern matching and information retrieval. Higher-level cognitive functions like information synthesis, inference and drafting still require human expertise.

In this paper, we investigate how GPT-4 performs at urban planning tasks that require higher levels of cognition. We study GPT-4's capability to interpret and respond to complex questions in the context of urban planning and design in Singapore. We describe the strengths of GPT-4 and highlight the prospects of application in the near future. We also discuss present capability gaps and provide

our outlook on how GPT technology is likely to impact the state of urban planning practice moving forwards.

Lines of Enquiry and Methodology

In the context of urban planning and design in Singapore, what are GPT-4's current abilities? What is its state of training and 'understanding' of Singapore-specific information, i.e., planning terms, definitions and guidelines? To what extent is there potential to augment planners' abilities with AI? If we look towards planners' potential use and application, are there red flags to be mindful of?

This exploration thus focuses on two broad areas, (i) planning knowledge, concepts and strategies and (ii) more technical-oriented data, analytics and tools. Five themes were explored - macro strategies and policies, the Master Plan and local planning, urban design, development control, and infrastructure planning.

Like ChatGPT, GPT-4 functions as a chatbot. Given that GPT-4 was trained on data prior to late-2021, conversation was first initiated with text from recent articles to test its ability to recognise the theme and its linkages to the deeper corpus of urban planning knowledge for each theme. Over 80 prompts ensued as a dialogue, covering the two focus areas across the five themes. As a quick exploration, no quantifiable or scientific methods were used to test GPT-4's abilities. The questions were phrased in a conversational manner with minimal technical jargon akin to an interview, to gauge GPT-4's responses relative to that of a human planner. The following sections highlight key observations.

Observations

Macro Planning Concepts & Strategies

The conversation on macro planning concepts and strategies was initiated with a newspaper commentary discussing meritocracy in Singapore, a broad social issue, with the intent to see how GPT-4 could connect such a topic with urban planning and policy-making.

GPT-4 gave well-formed textbook answers, demonstrating its ability to interpret the commentary theme, draw linkages to planning and policymaking and to lay out planning strategies (see example in Figure 1). It correctly covered various aspects ranging from land use zoning to spatial equity through accessibility and amenity provision. This ability could be related to the volume of available planning literature, especially from Western planning frameworks and academic sources. To a limited extent, it was able to relate to the Singapore planning context, referencing the roles and data sources of URA and other agencies.

Prompt: Analyse the text – How would this impact policies in urban planning and the built environment?
[inserted text from commentary]

GPT-4: The text highlights the importance of fostering an open and inclusive meritocracy in Singapore, while addressing four key priorities: resisting stratification, strengthening social safety nets, building a city with a preserved quality of life, and reinforcing Singaporeans' responsibilities towards one another. The analysis of this text suggests several potential impacts on urban planning and the built environment policies:

1. **Emphasis on diverse education and workplace opportunities:** Urban planning policies should focus on creating diverse and inclusive educational and work environments. This may involve designing mixed-use neighborhoods that support a variety of institutions, from schools and polytechnics to vocational training centers, promoting accessibility and opportunities for people with different skill sets and backgrounds.
2. **Enhanced social safety nets:** The focus on strengthening social safety nets may require the development of affordable housing, accessible public transportation, and inclusive public spaces to cater to the diverse needs of the population, particularly for lower-income workers and people with disabilities.
3. **City ecology and quality of life:** Urban planning policies should prioritize the preservation of the city's ecology and overall quality of life. This can be achieved through the development of green spaces, sustainable infrastructure, and measures to reduce pollution and environmental degradation. Additionally, planners should consider promoting walkability and mixed-use development to create vibrant and livable communities.
4. **Reinforcing social responsibility:** Encouraging social responsibility among Singaporeans requires fostering a sense of community and shared responsibility in the urban environment. Urban planning policies should consider the development of community spaces and facilities, such as community centers, parks, and recreational facilities, that promote social interaction, collaboration, and support among residents.
5. **Balancing meritocracy with inclusivity:** Urban planning policies should ensure that meritocracy does not inadvertently lead to stratification, by promoting an inclusive approach to development. This may involve

considering the needs of various socioeconomic groups, providing equal opportunities for all residents, and addressing potential barriers to social mobility in the design and implementation of urban policies.

Focus on future demands: In the context of workforce development, urban planning policies should be forward-looking and cater to the anticipated needs of the future. This may involve developing infrastructure and facilities that support emerging industries, technology, and innovation, as well as ensuring that the built environment is adaptable to accommodate future changes and demands.

Analysis: The four key priorities summarised from the commentary were generally accurate, and GPT-4 could link these well to urban planning concepts and strategies. However, the third summary point had a slightly different meaning, as the original text mentioned “continue to build a city whose ecology preserves its quality of life”, while GPT-4 shortened it to “building a city with a preserved quality of life”.

Figure 1 - Connecting a broad social issue with macro planning concepts & strategies

GPT-4 was also able to suggest relevant quantitative analytical methods related to the theme, for example, formulating a weighted approach to evaluate access to services and opportunities (see example in Figure 2). It developed the methodology further when prompted to consider different demographic profiles. GPT-4 could also provide technical code implementation of the methodologies in GIS tools such as ArcGIS Pro and QGIS. Probed further, it laid out a sound planning approach to address immediate versus longer-term changes and trends using scenarios, quantitative analysis and a monitoring framework. It could also engage in meaningful discussion on the need for planning practice to involve a balanced quantitative and qualitative approach.

Prompt: Suppose I have computed the spatial accessibility of residential dwelling units to a few key facilities such as parks, clinics, sports centres. How would I put these together into a composite index? What models or references have other cities used? Would you weigh one type of facility (e.g. clinic) more than another (e.g. cinema)?

GPT-4: Creating a composite index of spatial accessibility involves combining the accessibility measures for different types of facilities into a single metric. There are several ways to do this, and the choice of method and weights for different types of facilities may depend on the specific context, objectives, and preferences of the decision-makers.

One approach to creating a composite index is to use a weighted sum model. In this method, you assign a weight to each type of facility based on its relative importance in the context of your analysis. Then, you multiply the accessibility score for each facility type by its weight, and sum these weighted scores for each dwelling unit to get the composite index. For example, suppose you have the following accessibility scores for three types of facilities (parks, clinics, and sports centers) for each residential dwelling unit:

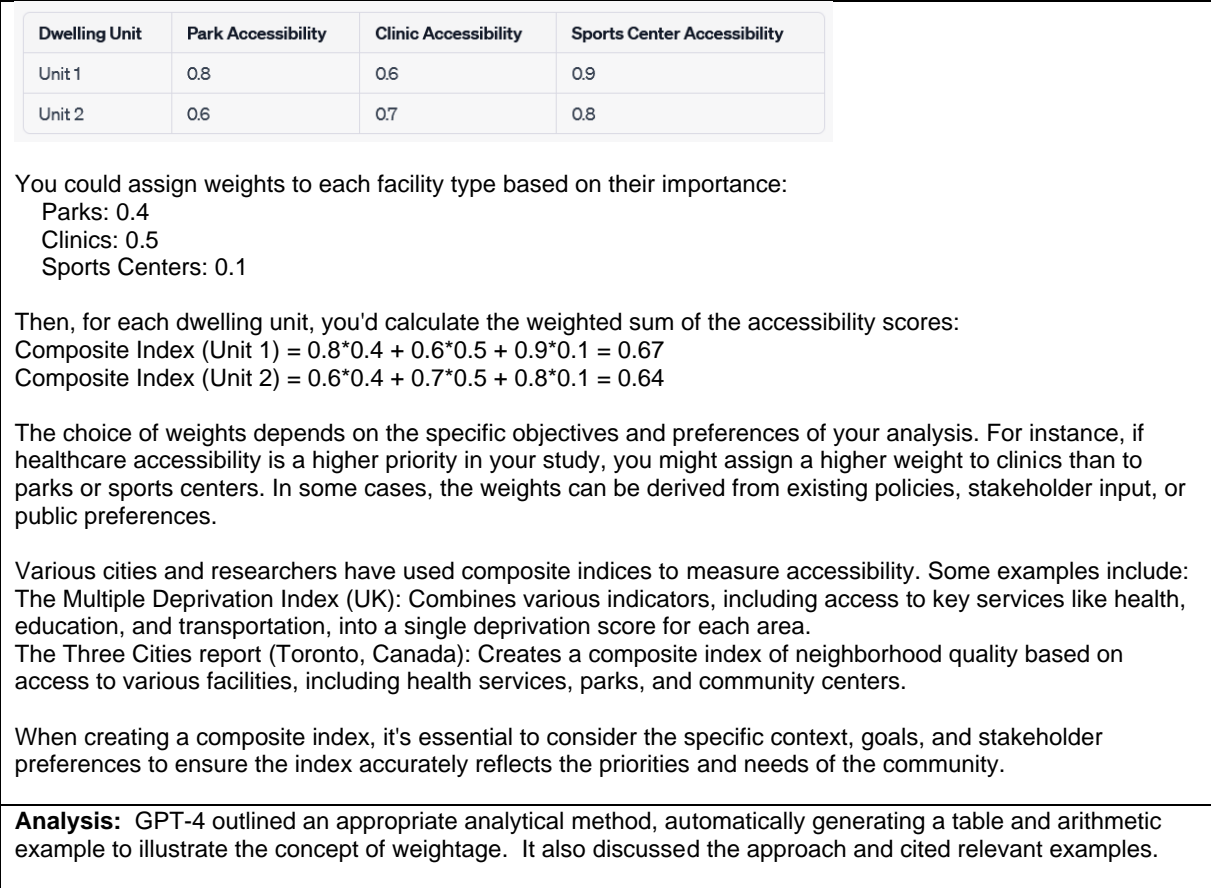


Figure 2 - Analytical method related to the planning theme

One of GPT-4's weaknesses was its inability to provide working links to web publications. This could be attributed to these documents being removed from source websites. Based on the conversation, GPT-4 operated at a level akin to an experienced urban planner, potentially exceeding human planners' ability to articulate concepts in an organised manner. Potential immediate applications include GPT-4 assisting planners to outline planning strategies, check for 'blind spots' in conceptual frameworks, draft concepts or talking points, outline analytical approaches, etc.

Master Plan and Local Planning

The conversation on the statutory Master Plan (MP) and local land use and planning considerations was initiated with text from a newspaper report relating to the en-bloc sale and potential redevelopment of a mixed-use development.

Like the preceding section, GPT-4 was able to provide sound textbook answers on more general and conceptual planning issues, such as mitigating local impacts arising from new development or redevelopment. When prompted, it could also respond with accurate information on a district's local heritage and character (see Figure 3). It can thus be useful for assisting planners in articulating concepts and issues, as well as surfacing general background information on local places.

Prompt: Are there specific considerations relating to the history or local character of Katong?
GPT-4: Yes, the Katong area has a rich history and distinct local character that should be taken into consideration during any redevelopment process. Some of these considerations include: Peranakan culture: Katong is known for its strong Peranakan heritage, which is a unique blend of Chinese, Malay, and Indonesian cultures. The area is home to many traditional shophouses, eateries, and shops offering Peranakan food, clothing, and handicrafts. Any redevelopment should respect and, if possible, integrate this cultural heritage to maintain the area's unique identity. Architectural heritage: The area features a mix of traditional shophouses, colonial bungalows, and modern residential buildings. Preserving and integrating elements of the architectural heritage into the new development can help maintain the local character and historical continuity. Urban identity: Katong has a vibrant and diverse urban identity, featuring a mix of commercial, residential, and recreational spaces. Ensuring that the new development complements and enhances the existing urban fabric and community amenities is essential for maintaining the area's distinct character. Local businesses: The area is known for its mix of traditional and modern businesses, including food establishments, retail shops, and service providers. Supporting and integrating these businesses into the new development can contribute to the local economy and maintain the area's commercial appeal. Community engagement: Engaging with the local community and incorporating their feedback and concerns into the planning process can help ensure that the new development meets the needs of residents and contributes to the overall well-being and development of the area. By considering these factors and working closely with local stakeholders, developers can help ensure that any redevelopment of the Katong Shopping Centre site respects the area's rich history and unique local character, while also providing modern amenities and contributing to the ongoing growth and development of the district.
Analysis: Whilst the information was not mentioned in the original text, GPT-4 had knowledge of the Katong area and responded with accurate information. This is likely from available web sources, such as articles on the history and character of Katong.

Figure 3 - Knowledge of local area

However, GPT-4 was unable to provide accurate responses relating to the MP land uses and planning processes such as rezoning, and it hallucinated in some instances (see example in Figure 4).

<p>Prompt: Does the Master Plan 2019 cater to uses such as funeral parlours or columbaria? What land use categories would these be under?</p>
<p>GPT-4: Yes, the Master Plan 2019 does cater to uses such as funeral parlors and columbaria. These facilities typically fall under the "Civic and Community Institution," "Community Use," or "Special Use" land use categories. These categories are designed to accommodate a diverse range of facilities and services that cater to the community's needs and requirements.</p> <p>Funeral parlors and columbaria serve essential functions related to end-of-life services and rituals, and they need to be carefully planned and integrated into the urban environment. The authorities must consider factors such as accessibility, the potential impact on surrounding land uses, cultural sensitivities, and environmental concerns when allocating land for these facilities.</p> <p>In Singapore, the Urban Redevelopment Authority (URA) ensures that appropriate sites are designated for funeral parlors and columbaria within the Master Plan, taking into account the factors mentioned above and the long-term needs of the community.</p>
<p>Analysis: Hallucination; the land use categories for funeral parlours and columbaria are "C&CI" and "Cemetery", respectively. These land use definitions are publicly available and have been in use well before 2021.</p>

Figure 4 - Knowledge of MP land use definitions

It was also unable to provide responses relating to geospatial data such as development names, addresses and corresponding land uses. At one point, when pressed for a data source, it wrongly suggested a non-existent but believable link, "datamall.ura.gov.sg". These shortcomings could be due to the gaps in its knowledge corpus, as well as functional limitations in its present form as an LLM with no geospatial capability. At this point in time, GPT-4 is not a reliable source of advice on the MP and other geospatial-related matters.

Urban Design

The conversation on urban design was initiated with text from a newspaper report relating to the success of precinct shopping malls, pivoting to questions relating to urban design concepts and strategies.

Again, GPT-4 was able to respond well on general urban design concepts and strategies, citing elements such as connectivity, diversity, public realm and streetscape design, etc. It was able to cite broad urban design strategies from various cities and even articulate relevant points on principle-based, prescriptive-based and performance-based approaches. However, as above, it was poor in its ability to provide working links to web publications and to quote accurate urban design guideline parameters such as building setbacks and heights.

When posed questions on scenarios with descriptions of physical characteristics (e.g., building storey height, street width), it was able to outline possible differences in spatial experiences between scenarios (e.g. sense of scale, visual interest) (see Figure 5). Through Scalable Vector Graphics (SVG) code, it could also provide simple sketches of the scenarios, as well as 3D visualisation in HTML/Three.js (see Figure 6).

Prompt: Alice stands in the middle of a 15m wide street. The street is flanked on both sides by shops which are 2 storeys in height. Betty stands in the middle of another street, which is 30m wide. The street is flanked on both sides by tall office skyscrapers which are 40 storeys in height. Discuss and compare Alice and Betty's spatial experiences.

GPT-4: Alice and Betty's spatial experiences would differ in several aspects, including the sense of scale, visual interest, natural light, and overall atmosphere. Let's discuss these factors in detail:

Sense of scale: Alice is standing in a 15m wide street with 2-storey shops on both sides, which creates a more intimate, human-scale environment. The low-rise buildings are closer to human proportions, which can make people feel more comfortable and connected to the built environment. In contrast, Betty is in a 30m wide street with 40-storey skyscrapers on both sides, which creates a more significant sense of scale. The towering buildings may make Betty feel small and insignificant in comparison, which can be intimidating or awe-inspiring, depending on personal preferences.

Visual interest: The 2-storey shops in Alice's environment likely offer a mix of architectural styles, colors, and signage that create visual interest and can engage pedestrians at the street level. The ground-level retail shops may have window displays, which can further contribute to the visual variety. In Betty's case, the tall office

skyscrapers may have less visual variety at the street level, as their primary function is to provide office space, and their design may be more repetitive and uniform.

Natural light: The narrower street in Alice's environment could lead to less sunlight reaching the street, depending on the orientation and time of day. In some cases, this may result in a cozier atmosphere. On the other hand, Betty's wider street is likely to receive more sunlight due to the increased distance between the buildings. However, the tall skyscrapers might also cast long shadows on the street, which can create a cooler and darker environment, especially during the morning and late afternoon hours.

Overall atmosphere: The low-rise, mixed-use environment in Alice's case is more likely to create a lively, pedestrian-friendly atmosphere. The presence of shops, cafes, and other small businesses may encourage walking, socializing, and street-level activity. In contrast, the high-rise office environment in Betty's case might feel more sterile, formal, and less inviting for pedestrians. The street life may be more business-oriented, with limited opportunities for leisure and social interaction.

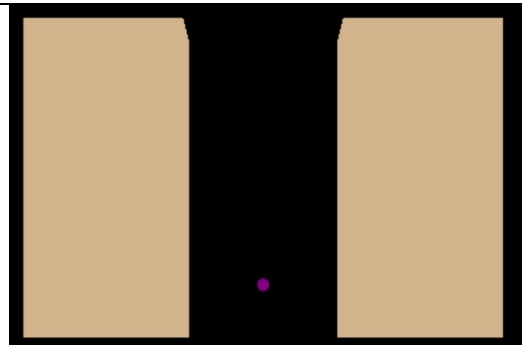
In summary, Alice's environment provides a more intimate, human-scale, and potentially more visually engaging experience, while Betty's environment offers a grander sense of scale and may feel more formal and business-oriented. Each person's preference for such environments would depend on individual tastes, comfort levels, and the purpose of their visit.

Analysis: GPT-4 inferred from the description and responded with relevant points on spatial experience, including sense of scale, visual interest and overall atmosphere. However, it wasn't clear whether it considered and compared the different building height-street width ratio, where the buildings along Alice's street are lower in height, even though the width is narrower.

Figure 5 - Understanding of space and human experience

Prompt: Generate Javascript code to illustrate Alice's street in 3D.

GPT-4: Creating a 3D visualization in JavaScript typically involves using WebGL and a library such as Three.js. Here's a simple example of JavaScript code using Three.js to create a 3D visualization of Alice's street, with Alice in the middle and 2-storey buildings on both sides. This code sets up a basic 3D scene using Three.js with a camera, renderer, and simple geometries to represent Alice and the buildings. The scene is rendered and animated using requestAnimationFrame. Please note that this is a very basic example, and you may need to adjust the parameters and settings for a more accurate and refined visualization.



[Sketch generated from HTML / Three.js code]

Prompt: For the above 3d visualisation, regenerate the code to set the camera at street level. Give some colour to the street and to the background sky.

GPT-4: I can provide a Three.js example, a popular 3D library for WebGL, to create a 3D visualization of Alice's street with the camera at street level and added colors for the street and sky. First, include the Three.js library in your HTML file: [Three.js code generated]. Then, add the following JavaScript code to create the 3D scene:



[Sketch generated from HTML / Three.js code]

Analysis: GPT-4 set the same dimensions for the buildings as similar as a preceding 2D sketch, assuming a 5m height for the 2-storey building. For the first 3D sketch, as the camera angle wasn't specified, GPT-4 set the camera angle viewing downwards. Subsequently, as prompted, GPT-4 created a ground plane, gave it a different colour from the buildings, and added a blue sky background. It also set the camera at street level, at Alice's 1.7m height specified in a previous prompt.

Figure 6 - Generating code for 3D space

In a stretched test of its spatial abilities, it was able to illustrate in SVG a simple site layout based on textual descriptions and rules (e.g., setback area, distance between building footprints) (see Figure 7). However, despite several attempts in refining the prompt, it did not go beyond simple layout and stopped short of obeying the instruction to fill the simulated parcel. GPT-4 instead responded, very astutely, that as a LLM it was unable to implement a packing algorithm.

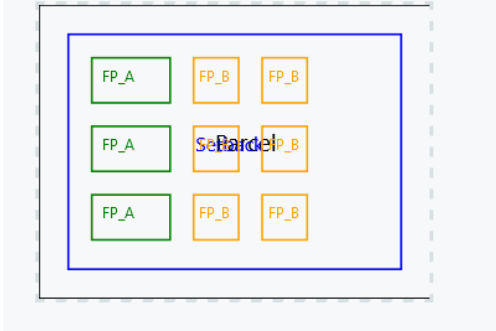
<p>Prompt: Generate SVG code to plot the following. (a) 200px x 150px rectangle labelled as 'Parcel'. (b) An internal offset of 15px from the rectangle to represent building setback, labelled as 'Setback'. (c) Small rectangles measuring 40px x 23px labelled as 'FP_A'. (d) Small squares measuring 23px x 23px labelled as 'FP_B'. (e) Fit in as many 'FP_A' and 'FP_B' as possible within 'Parcel', to maximise the total area of 'FP_A' and 'FP_B'. (f) No part of the 'FP_A' or 'FP_B' shall be located within 'Setback'. They shall not even cross the 'Setback' line. (g) Measured from the edge, there shall be a minimum 12px distance separating all 'FP_A' and 'FP_B'.</p>	
<p>GPT-4: Here's the SVG code for the requested illustration. It consists of a 200px x 150px 'Parcel' rectangle, a 170px x 120px 'Setback' rectangle, and multiple 40px x 23px 'FP_A' rectangles and 23px x 23px 'FP_B' squares placed within the 'Parcel' to maximize the total area, without overlapping the 'Setback' and maintaining a minimum 12px distance between each other.</p>	 <p>[Sketch generated from GPT-4's SVG code]</p>
<p>Prompt: There is still empty space to fill more shapes.</p>	
<p>GPT-4: You're right, there is more space available. However, creating an optimal packing algorithm is a complex task, and it's outside the scope of this AI model. The example I provided demonstrates the concept, but it is not an optimal solution. If you are interested in finding an optimal solution, you might consider researching packing algorithms or using specialized software or libraries that are designed to solve packing problems.</p>	
<p>Analysis: GPT-4 is able to understand the instructions well, setting out, labelling and colour-coding the different elements, and translating these to SVG code. In trying to push it to fill the space, GPT-4 seemed to recognise the intent of the prompt and responded very astutely that as a LLM, it was unable to implement a packing algorithm!</p>	

Figure 7 - Generating code for spatial plan layout

From tests of its image recognition abilities through web links to photos, GPT-4 hallucinated in several areas, wrongly attributing the Singapore skyline to Hong Kong in one case and mentioning several elements which were not present in the

image in another case. One possibility was that it had formulated its response based on the metadata of the image or the content from the image's website. At this point in time, GPT-4 has limitations in image recognition.

Development Control

The conversation on development control was initiated with text from a regulatory circular pertaining to definitions of floor area, to test GPT-4's ability relating to Singapore-specific regulations and technical definitions.

GPT-4 was able to respond well on general development control concepts. However, given the highly contextual and specific nature of guidelines, definitions, parameters and regulatory processes, GPT-4 was poor in its ability to provide accurate responses in several areas. Some of its responses appeared to be formulated based on information or practices from other cities. In general, GPT-4 is currently unable to meaningfully assist professional qualified persons (QPs) or planners on regulatory matters and processes (see example in Figure 8).

Prompt: How about car park areas? Are vehicle parking areas and driveways computed as GFA?
GPT-4: In Singapore, car park areas are treated differently in GFA calculations depending on their classification. According to the Urban Redevelopment Authority (URA) guidelines: Open surface car parks and car park areas on rooftops without any covering are not computed as GFA. Covered car park spaces, such as those within multi-story car parks, basements, or under a podium, are included in GFA calculations. Driveways, on the other hand, are generally not considered as part of GFA calculations. Driveways serve as access routes for vehicles to enter and exit a development, and they do not contribute to the building's floor area, as they are not enclosed spaces. It's essential to be aware of the specific guidelines and criteria set by the URA for car park and driveway areas, as these may vary depending on the type of development and location.
Analysis: Hallucination. Carpark and driveway areas are not computed as GFA. GPT-4 did caveat that specific guidelines may vary.

Figure 8 - Understanding of development control regulations

GPT-4 could respond well on some technical issues, such as explaining the Industry Foundation Classes (IFC) format used in regulatory submissions and could

even describe snippets of IFC code. However, due to GPT-4's limited input length, the code was truncated and could only represent a very small portion of a development.

Infrastructure Planning

The conversation on infrastructure planning was initiated with text from a newspaper report on an upcoming major road infrastructure.

GPT-4 was able to respond well on general infrastructure planning concepts, including the interaction between infrastructure and land use planning. For example, it provided well-formed responses on key processes and considerations in infrastructure planning (e.g., identifying needs, data collection and analyses, conceptual planning and design, stakeholder engagements, impact assessments, detailed design & engineering, regulatory approvals, financing & procurement, construction & monitoring, etc.). As above, its responses appeared to be formulated based on information or practices from other cities.

When prompted on transport planning, it responded accurately with different analytical methods, outlining strengths and weaknesses, as well as some commonly-used tools. It was also able to illustrate a simple example of trip generation, based on the traditional four-step transportation demand modelling method, where it assumed rates from published sources and performed simple calculations (see Figure 9).

Prompt: Assume a new mixed-use development with: (a) 17,000 sqm residential use, or 200 dwelling units, (b) 10,000 sqm commercial office use, (c) 3,000 sqm commercial retail, food & beverage use. Based on these uses, compute an estimated trip generation, explaining your assumptions.

GPT-4: To compute the estimated trip generation for the given mixed-use development, we need to make some assumptions about the average trip rates for each land use type. Trip rates can vary significantly depending on factors such as location, local travel patterns, and demographic characteristics, so it's essential to consult local data sources or previous studies to determine appropriate trip rates for your specific context. For the purpose of this example, we'll use the following assumptions based on generic trip rates from the Institute of Transportation Engineers (ITE) Trip Generation Manual:

Residential (200 dwelling units): Average trip rate: 10 trips per dwelling unit per day. This includes both inbound and outbound trips.

Commercial office (10,000 sqm): Average trip rate: 1.5 trips per 100 sqm per day. This includes both inbound and outbound trips.

Commercial retail, food & beverage (3,000 sqm): Average trip rate: 40 trips per 100 sqm per day. This includes both inbound and outbound trips.

Now, we'll compute the estimated trip generation for each land use type:

Residential trip generation:

200 dwelling units * 10 trips per dwelling unit = 2,000 trips per day

Commercial office trip generation:

(10,000 sqm / 100 sqm) * 1.5 trips per 100 sqm = 150 trips per day

Commercial retail, food & beverage trip generation:

3,000 sqm / 100 sqm) * 40 trips per 100 sqm = 1,200 trips per day

Total estimated trip generation for the mixed-use development:

2,000 (residential) + 150 (office) + 1,200 (retail, food & beverage) = 3,350 trips per day

Keep in mind that these trip generation estimates are based on generic trip rates and should be considered as a starting point for further analysis. It's essential to consider local context, transportation infrastructure, and travel patterns to refine these estimates and develop more accurate transportation demand forecasts for the mixed-use development.

Analysis: The method is sound. Based on the ITE manual, trip rate of 10 applies for single-family detached housing. High-rise apartments will be lower at 5.2, but the type of dwelling wasn't specified in the prompt. Trip rates will be different for the Singapore context, but this wasn't specified in the prompt. Choosing the trip generation rate typically requires some interpretation on the land use, including the overall scale and quantum of the land use components.

Figure 9 - Transportation demand modelling concepts

When stretch-tested on its spatial abilities, it faced challenges in deconflicting two underground pipes in relation to their depths and diameters. Similar to the urban design example, it however was able to generate a 3D sketch of an underground pipe based on textual descriptions.

Discussion and Conclusion

GPT-4 is strong in its ability to articulate concepts and strategies relating to urban planning, policies, urban design, etc., providing textbook responses in a very well-organised manner. Its ability to argue, for example, outlining strengths and weaknesses of different approaches, is structured and well-formed, akin to an experienced urban planner's ability.

GPT-4 is also strong in suggesting methodologies of technical analyses, including providing detailed steps or code for running GIS software (both ArcGIS and QGIS). Based on textual descriptions, it can also generate code to illustrate simple diagrams both in 2D and 3D, e.g., sketch section of a street.

GPT-4 is weak in providing internet reference links, many of which did not work. As a LLM, it does not have ability to extract or present geospatial data. It is also weak in responding to Singapore-specific domain knowledge such as technical land use descriptions, urban design / development control parameters (e.g., setback distances), and hallucinated in several instances. It did not perform well in image recognition tasks.

GPT-4's abilities are likely related to its training data. If the mechanisms in GPT-4 are functionally similar to simpler machine learning models, then the content in its knowledge corpus will likely determine the quality of its response. This points to the need to better understand the source of its knowledge, and the scope of supplementary information it should be provided with to function effectively in specialised domain applications, in this case, urban planning.

The speed at which GPT-4 and other LLMs is developing is astonishing. There is already potential in the immediate-term to deploy GPT-4 to assist planners in its areas of strength, such as articulating concepts and strategies, and outlining technical methodologies. If incorporated into planning processes and workflows, it can help planners perform more well-rounded analyses and reap productivity gains on daily tasks. However, planners must be well-informed on GPT-4's areas of weaknesses, in particular where hallucinations appear very well-phrased and hence may be difficult to detect.

An AI tool like GPT-4 does not take away the need to continue building urban planners' competencies, knowledge and understanding of cross-domain issues, as well as their use of data, tools and analytics. The onus, at least at this point in time, is still on human planners to fact check, evaluate logical flows, reason, think critically and apply knowledge while working alongside these tools.

References

- Ayeni, B. (2003). The design of spatial decision support systems in urban and regional planning. In *Decision support systems in urban planning* (pp. 20-33). Routledge.
- Batty, M. (1976). *Urban modelling* (pp. 20-48). Cambridge: Cambridge University Press.
- Bubeck, S. (2023). Sparks of Artificial General Intelligence: Early experiments with GPT-4. arXiv.org. <https://arxiv.org/abs/2303.12712>
- Daniel, C. (2023). ChatGPT: Implications for Planning. American Planning Association. <https://www.planning.org/publications/document/9268026/>
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning, *nature* 521(7553) (pp. 436-444).
- OpenAI. (2023). GPT-4 Technical Report. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2303.08774>