

JAN 2024

The Potential and Limitations of AI in the Site Selection Industry



The Potential and Limitations of AI in the Site Selection Industry



About the Site Selectors Guild

The Site Selectors Guild is the only association of the world's foremost professional site selection consultants. Guild members provide location strategy to corporations across the globe and for every industry, sector, and function. Founded in 2010, the Site Selectors Guild is dedicated to advancing the profession of international corporate site selection by promoting integrity, objectivity, and professional development. Members are peer-nominated, vetted, and must demonstrate significant professional location advisory experience. In 2022, Guild members facilitated the announced creation of approximately 120,000 jobs and \$268 billion in announced capital investment by corporate clients, which required location analysis across 81 countries. Guild membership is the highest standard in the site selection industry.

Table of Contents

Key Findings	06
What is AI and How is it Currently Utilized by Site Selectors?	07
Putting AI to the Test	09
Case Study #1: A Software Company Relocating from San Francisco	10
Case Study #2: A Beverage Manufacturing Bottling Facility	14
Overall Performance	17
Conclusion	19

Copyright © 2024 Site Selectors Guild. All rights reserved.

No part of this document may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without permission in writing from the publisher, except by reviewers, who may quote brief passages in a review.

Guild Members

Angelos Angelou, *AngelouEconomics*
 Alex Ash, *Alpha Strategies*
 Sydney August, *CBRE, Inc.*
 Olaf Babinet, *Deloitte Consulting GmbH*
 Johan Beukema, *BCI Global*
 Robert Boehringer, *KPMG*
 Tracey Hyatt Bosman, *Biggins Lacy Shapiro & Company*
 Rene Buck, *BCI Global*
 Darin Buelow, *Deloitte Consulting LLP*
 Didi Caldwell, *Global Location Strategies*
 Ron Crum, *CSRS, Inc.*
 David Dera, *Newmark*
 Jan Desmaretz
 Dennis Donovan, *Wadley Donovan Gutshaw Consulting*
 Andreas Dressler, *Location Decisions*
 Courtney Dunbar, *Burns & McDonnell*
 Tess Fay, *Global Location Strategies*
 Jay Garner, *Garner Economics, LLC*
 Larry Gigerich, *Ginovus*
 Jeannette Goldsmith, *Strategic Development Group*
 Ford Graham, *McGuireWoods Consulting, LLC*
 Taylor Gravois, *CSRS, Inc.*
 William Hearn, *Site Dynamics LLC*
 Robert Hess, *Newmark*
 Jason Hickey, *Hickey & Associates*
 Matt Highfield, *Deloitte Consulting LLP*
 Bradley Lindquist, *Newmark*
 John Longshore, *Newmark*
 Natalie Lynch, *Deloitte Consulting*
 Dennis J. Meseroll, *Tractus Asia Limited*
 Denise Mullis, *J.M. Mullis, Inc.*
 Mike Mullis, *J.M. Mullis, Inc.*
 Tara Nicholson, *Deloitte Consulting LLP*
 Sarah Raehl, *Deloitte Consulting*
 Andrew Ratchford, *Site Selection Group*
 Scott Redabaugh, *JLL Business Consulting*
 James Renzas, *BCI Global*
 Wolfgang Riedel, *Riedel Consulting*
 John Rocca, *JLL*
 Matt Ryder, *Avison Young*
 Don Schjeldahl, *DSG Advisors*
 Ulrich Schmidt, *KPMG*
 Phil Schneider, *Schneider Strategy Consulting, LLC*
 Spencer Schobert, *Newmark*
 Alexandra Segers, *Tochi Advisors LLC*

Andrew Shapiro, *Biggins Lacy Shapiro & Company*
 Tom Stringer, *BDO*
 Chad Sweeney, *Ginovus*
 Gray Swoope, *VisionFirst Advisors*
 Jerry Szatan, *Szatan & Associates*
 Matt Szuhanj
 Rajeev Thakur, *JLL*
 Monty Turner, *Colliers International*
 Elias van Herwaarden, *Colliers International*
 Patsy Van Hove, *IBM-PLI (Plant Location International)*
 Leslie Wagner, *Ginovus*
 Gregg Wassmansdorf, *Newmark*
 Matthijs Weeink, *JLL*
 Sarah White, *Global Location Strategies*
 Mark Williams, *Strategic Development Group, Inc.*
 Kim Williams Davis, *Quest Site Solutions*
 Carter Wood, *Ernst & Young, LLP*
 Gary Yates, *Stream Realty Partners*

Board Chair

Larry Gigerich, *Ginovus*

Vice Chair

Seth Martindale, *CBRE, Inc.*

Board Treasurer

Philip Schneider, *Schneider Strategy Consulting*

Board Secretary

John Rocca, *JLL*

Board Members

Michelle Comerford, *BLS & Co.*
 Chris Lloyd, *McGuireWoods Consulting, LLC*
 Gregg Wassmansdorf, *Newmark*

Board Advisors

Kim Moore, *Newmark*
 Larry Moretti, *LFM Corporate Location Solutions*

Research Collaborators

- Development Counsellors International (DCI)
- Seth Martindale, *Senior Managing Director, CBRE, Inc.* and Gregg Wassmansdorf, *Senior Managing Director, Newmark*
- Dr. David Chiu, *Professor of Mathematics and Computer Science, University of Puget Sound*

The Potential and Limitations of AI in the Site Selection Industry

No industry sector – the site selection industry included – is immune from the disruption brought on by the rise of artificial intelligence (AI). The potential benefits are numerous, including greater efficiencies through the automation of tasks, enhanced decision-making, and motivating an innovation mindset. With these opportunities, it's a common sentiment that if an industry is not already using AI, it is already behind.

Yet, as the saying goes, “With great power comes great responsibility,” and AI has unleashed immeasurable new powers. If not harnessed or managed properly, AI could result in outcomes that are, at best, misdirected and, at worst, financially or operationally damaging to an enterprise or organization.

Representing an industry that manages high-stakes location decision-making for companies across the globe, the Site Selectors Guild – the only association of the world's foremost professional site selection consultants – recognizes the importance of understanding how AI tools might be successfully leveraged without sacrificing the quality and accuracy of site location searches.

To take some initial steps toward understanding the potential impacts AI could have on the site selection industry, the Guild undertook two research projects in the fall of 2023.

- A **survey** of members of the Site Selectors Guild was conducted to gauge their current and anticipated use of AI in the site selection process, the greatest opportunities, and existing limitations to provide an overview of the integration of AI into the site selection industry.
- A **research study** designed to “test” three AI platforms (ChatGPT 3.5, ChatGPT 4.0 and Bard) against the human-generated outcomes of two real-world site selection projects – one office and one industrial. An independent expert in mathematics and computer science with no previous knowledge of site selection or the case study location decisions managed the project.




Key Findings

The independent research study outcomes supported the overall conclusion from the survey of Guild members that artificial intelligence has the potential to introduce efficiencies into the site selection process and free human resources to focus on more complex project elements, but in AI's current state, all platforms would require significant human oversight and vast differences exist in the outcomes generated by the different platforms.

- In the case of the industrial project, **there was almost no overlap between the AI-generated shortlists of metro location recommendations and the consultant-generated shortlist for the real project.** Through a narrative query in ChatGPT 4.0 (where the inputs are entered as one large paragraph), only one of the locations on the human-generated shortlist was also identified by the platform as a potential location for the project. This was also the case using the Bard platform.
- In the case of the office project, the results were more promising, specifically using the narrative approach. **The ChatGPT 3.5 platform identified nine of the 17 potential office metro locations** that also appeared on the consultant-generated list, followed by ChatGPT 4.0 results, which identified seven of the locations on the consultant list, and finally, Bard, which identified three results that were also on the consultant list.
- In the least successful outcome, a piecemeal query was used whereby the project search criteria inputs are entered one descriptive sentence at a time. When the industrial project criteria were entered into ChatGPT 3.5, and despite repeated attempts, **the platform was unable to generate any recommended industrial project locations.** Instead, the free version of ChatGPT generated the following response:

“ To create a ranked list of cities that meet your criteria, **I recommend working with a professional site selection consultant or conducting thorough research and analysis** using location-specific data. These experts can help you assess each city's suitability based on your project's unique needs and priorities.” – *ChatGPT 3.5*

Additional details on the Guild survey and the research study, where there is opportunity and where limitations persist for AI, are provided on the following pages.




What is AI and How is it Currently Utilized by Site Selectors?

Artificial Intelligence (AI) is the capability of computer systems or algorithms to imitate intelligent human behavior. It was in the 1950s that key advances were made in artificial neural networks, early attempts to measure machine intelligence, and the origination of the term “artificial intelligence”. Recent advances in computational power and breakthroughs in both mathematical and computer sciences have unleashed a torrent of new AI products and applications.


According to the survey of Guild members in November 2023, a notable percentage already utilize AI platforms for various day-to-day functions. When asked, “Do you use AI for office and administrative functions?” 40% of Guild members report using AI for these functions.

Figure 1. Guild Members Reporting Use of AI for Office and Administrative Functions



Of those Guild members who have used AI for office or administrative functions, 84% agree or highly agree with the statement, “My firm is using AI more to improve the efficiency of our office administrative functions and to increase productivity than in the site selection process or analysis.” This indicates a general unwillingness to use AI for functions directly related to the analysis and final decision-making in a client project. While 38% of Guild members report using AI in some capacity for a client’s project, there was strong consensus on the limitations of using AI in the actual location search.

Figure 2. Guild Members Reporting Use of AI (i.e., ChatGPT, Jasper, Scribe, etc.) to Assist with a Client’s Location Search



Of those reporting they have previously used AI to assist with a client’s location search, nearly all report that the primary use was to distill and analyze the tremendous amount of secondary data required for a location search. Other uses include reviews of state legislation and permitting requirements or copy-editing.


“ It is another tool that we use to help our analysts evaluate the tremendous amount of secondary data that is available in the public domain. However, **it is important to note that it is only one tool and not the only one we use.** We also use it to assist in the copy editing of our reports. AI is currently limited in its usefulness to the evaluation of publicly available data.”

– *Dennis J. Meseroll, Co-Founder and Executive Director, Tractus Asia Limited*

Sixty-one percent of respondents reported that using AI added value to the process by freeing human resources to assist with other project elements. Of those who did not see any added value after using AI, the primary critique was the time and specificity required with the original inputting of data to ensure valid results.

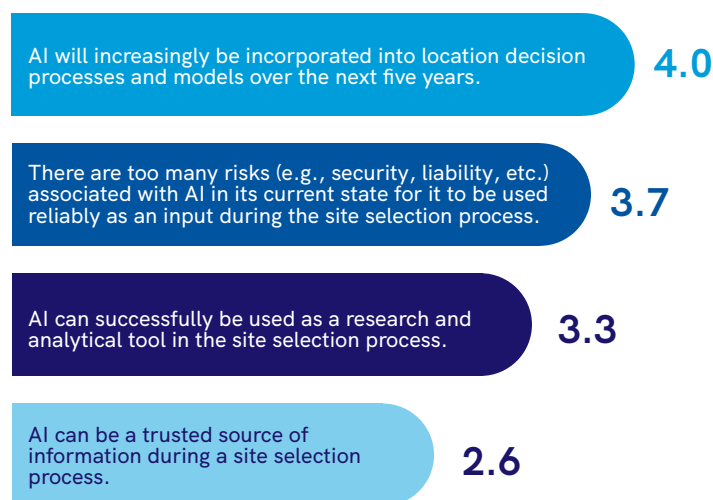
There are also mixed reactions on how trustworthy AI-generated information is, evidenced when Guild members were asked to provide their level of agreement with the following statement: “I trust the results generated by the generative AI platform used during our most recent site selection project.” With the majority of consultants using AI for administrative or processing purposes, still only 33% report that they trust the results generated by the platform.

Figure 3. Level of Agreement with the Statement “I trust the results generated by the generative AI platform used during our most recent site selection project.”
(1=don’t agree at all, 5=highly agree)



The overall sentiment among Guild members is that AI could potentially be a time-saving tool that allows more expedited analyses; however, in its current form, AI-generated results need to be treated with a high degree of caution requiring skilled, professional scrutiny for it to be a trustworthy input to the site selection process.

Figure 4. Level of Agreement with Select Statements
(1=don’t agree at all, 5=highly agree)



Guild members understand that artificial intelligence could be a powerful tool that can potentially alter – and ultimately support – the industry, but it also requires close oversight to be managed and utilized effectively.

Putting AI to the Test

To more fully understand the opportunities and limitations of AI as a tool in the location decision-making process, an independent study was commissioned by the Guild and managed by Development Counsellors International (DCI) for an impartial “test” that compared the outcomes of two real-world, previously conducted site selection projects against the outcomes generated by publicly available AI tools.

Conducted over a three-week period in October 2023, the AI site selection study included the following:

- DCI independently identified and commissioned an academic consultant (a doctored professor of computer science) at the University of Puget Sound to conduct the analysis.
- Members of the Site Selectors Guild provided the parameters for two recently completed site location searches:
 - A software company’s headquarters being relocated from San Francisco.
 - A manufacturing company locating a new facility to increase capacity.
- The two confidential projects were selected to represent two different sectors with distinct location and workforce criteria, dissimilar project specifications with each including more than a dozen variables, and unique sensitivities to property market dynamics and incentive structures.
- The academic consultant was not privy to the final shortlist of locations for either project before entering project information and data into various online AI platforms and developing an independent list of recommended locations.

Three generative AI platforms were tested: Google Bard, ChatGPT 3.5, and ChatGPT 4.0. These are arguably the top publicly available generative AI platforms and are most likely to be utilized by a company self-performing a site selection project. For each project example, queries were curated and prepared for input into each platform by the independent academic consultant to incrementally navigate and strengthen the contextual model of the conversation with each AI tool. Once all the site



parameters had been input, a series of iterative requests were made to obtain a list of 20 locations, ranked by suitability to the company’s needs. This was followed up with a request for 20 cities that the chatbots rejected.

Certain “best practices” for interacting with chatbots were employed. For instance, prompts were written to set the proper context, meaning the AI platforms were asked to problem-solve from the point of view of a “company executive” or a “company strategist.”

For this project, the chatbot had to be “coaxed” to accept multiple input parameters. To this end, the queries were formatted in two ways: (1) narrative form and (2) piecemeal. The narrative form combines the queries into a single large paragraph, and is entered as one input. In piecemeal form, each sentence of the query was input separately, requiring the chatbots to respond in a stepwise or iterative process. Interestingly, each input format elicited significantly different results.

Case Study #1: A Software Company Relocating from San Francisco

The first case study tested was for a corporate headquarters relocation of a software company moving from San Francisco, citing heightened costs and talent-recruitment challenges. The problem statement, criteria inputs for the exercise, and follow-up questions included the following.

1. From the point of view of a San Francisco software company strategist, identify a new metropolitan region to relocate your corporate headquarters. The expected budget for this relocation is \$10,000,000. The new location will house under 300 employees in a space with 50,000 square feet. Downtown or suburban areas can be considered if they have a robust public transportation system. Direct, nonstop flights to San Francisco are required from this new location. The ideal new headquarters must have the following: (1) a business-friendly operating environment with substantial economic incentives from state and/or local governments; (2) B-1 allows you to hire software-development talent at lower costs and attracts talent both globally and domestically; (3) provides a better quality of life and lower cost of living for your employees; and (4) has access to universities with strong computer science programs to cultivate new product research and workforce training and recruitment. Rank 20 metro areas according to these desirables and justify the ranking of each metropolitan area that is listed.
2. There needs to be at least three move-in ready facilities per metro area that you list.
3. List 20 cities that you rejected, and explain why.
4. Can you give me the raw data you used to rationalize <<INSERT CITY>> as the top choice?



Each AI system was asked to provide a ranked list of the top 20 metro areas and to justify the ranking of each metropolitan area that is listed, as well as 20 metro areas that were rejected.

Based on the inputs, the following results were generated across the three AI platforms using the two different inquiry methods. Only the top 17 results are shown to align with the human-generated list. Table 1 shows the shortlists generated when the parameters were entered into the three platforms in a piecemeal format compared to the consultant-generated outcomes. As seen, the results are mixed.

- Considering the 17 locations on the consultant-generated list, the most overlap occurred using the Bard platform which generated seven of the same locations.
- ChatGPT 4.0 generated six locations that also appeared on the consultant's list.
- There was the least overlap using the ChatGPT 3.5 platform which only identified five locations in the top 17 that also appeared on the consultant list.
- **Six locations on the consultant's list did not appear in any of the three AI platforms tested.**

It is also important to note that in most consultant searches, there will be three finalist locations in which site visits will occur. Combined, **the three platforms generated six different locations as the top three ranked locations.**

Table 1. Office Shortlist Results Generated by the AI Platforms (Piecemeal Approach) versus the Consultant-Generated Shortlist

RANK	ChatGPT 3.5 (piecemeal)	ChatGPT 4.0 (piecemeal)	Bard (piecemeal)	Human Generated
1	Austin	Atlanta	Austin	Atlanta
2	Boston	Dallas	Raleigh-Durham	Dallas
3	Raleigh-Durham	Raleigh-Durham	Denver	Charlotte
4	Salt Lake City	Austin	Nashville	Albuquerque
5	Nashville	Seattle	Salt Lake City	San Antonio
6	Phoenix	Washington, D.C.	Pittsburgh	Denver
7	Denver	Denver	Charlotte	Orlando
8	Seattle	Phoenix	Atlanta	Phoenix
9	Dallas	Nashville	San Antonio	Kansas City
10	San Diego	Miami	Kansas City	New Orleans
11	Washington, D.C.	Charlotte	Minneapolis	Boise
12	Atlanta	Tampa	Seattle	Nashville
13	Chicago	San Diego	San Francisco Bay Area	Columbus
14	Minneapolis-St. Paul	Minneapolis-St. Paul	Boulder	Tucson
15	Portland	Salt Lake City	Columbus	Cincinnati
16	Philadelphia	Portland	Indianapolis	Indianapolis
17	Houston	Philadelphia	Jacksonville	Louisville

Vastly different results were seen using a narrative approach (combining the queries into a single large paragraph, which is entered as one input). The success of the platforms was exactly the opposite of what was seen in the piecemeal approach.

- Considering the 17 locations on the consultant-generated list, the Bard platform was the least successful and only generated three locations on the consultant list.
- ChatGPT 4.0 generated seven locations that also appeared on the consultant’s list.
- There was the most overlap using the ChatGPT 3.5 platform which identified nine locations in the top 17 that also appeared on the consultant list.
- **Eight metros were on the consultant list that did not appear in any of the AI platform lists.**

Based on the results, the narrative approach generated the most overlap between all of the AI platforms and the consultant list. However, **the top-ranked location across all platforms (Austin) did not appear at all on the consultant-generated list** and the third-ranked location on the consultant list only appeared in the ChatGPT 3.5 generated list and ranked #15.

Table 2. Office Shortlist Results Generated by the AI Platforms (Narrative Approach) versus the Consultant-Generated Shortlist

RANK	ChatGPT 3.5 (narrative)	ChatGPT 4.0 (narrative)	Bard (narrative)	Human Generated
1	Austin	Austin	Austin	Atlanta
2	Raleigh-Durham	Atlanta	Dallas	Dallas
3	Seattle	Raleigh-Durham	Atlanta, GA (suburban)	Charlotte
4	Denver	Denver	Raleigh-Durham	Albuquerque
5	Atlanta	Pittsburgh	Denver	San Antonio
6	Nashville	Dallas	Seattle	Denver
7	Phoenix	Seattle	Boston	Orlando
8	Dallas	Salt Lake City	Atlanta	Phoenix
9	Minneapolis-St. Paul	Minneapolis-St. Paul	Washington, D.C.	Kansas City
10	Salt Lake City	Portland	Austin (suburban)	New Orleans
11	Portland	Tampa-St. Petersburg	Dallas, TX (suburban)	Boise
12	Tampa	Phoenix	Atlanta, GA (suburban)	Nashville
13	San Diego	Kansas City	Raleigh-Durham-Chapel Hill, NC (suburban)	Columbus
14	Columbus	Indianapolis	Denver-Aurora, CO (suburban)	Tucson
15	Charlotte	Miami	Seattle-Tacoma-Bellevue, WA (suburban)	Cincinnati
16	Indianapolis	Nashville	Minneapolis-St. Paul-Bloomington, MN (suburban)	Indianapolis
17	Kansas City	Columbus	Boston, MA (suburban)	Louisville

There is clearly some overlap between the human-generated and the AI-generated lists. However, **there is tremendous variation even among the AI-generated lists and depending on how the queries are entered (i.e., piecemeal versus narrative)**. In the case of the narrative approach, there were eight locations generated by the consultant that did not appear on any of the AI-generated lists, which raises some red flags, particularly in light of the fact that all of the platforms identified the same top-ranked location, but that location did not appear at all on the consultant-generated list. Given the high-level criteria and narrowly defined objectives of the search, one might expect there would only be a few U.S.-based metros that could meet the criteria, so the overlap that does occur is not surprising but the misalignment between the other locations identified by the consultant and the platforms does raise concerns.




Case Study #2: A Beverage Manufacturing Bottling Facility

The second case study required a new location for a beverage manufacturing facility, and this project's complexity presented clear challenges to the bots. While ChatGPT 4.0 and Bard produced seemingly "better" results, the chatbots cannot currently produce anything resembling a human-generated location recommendation list.

The problem statement, criteria inputs, and follow-up questions included the following.

1. From an executive's perspective for a large industrial bottling (beverage and food) company, identify potential locations for a new manufacturing building. This new building must be 600,000 square feet on 35 to 40 acres of land with industrial zoning or rezoning capabilities. The location should be close enough to large cities and population centers but must avoid proximity to residential areas due to the company's heavy manufacturing process. Employees are expected to drive to the new location, so access to high-volume transportation is not required.
2. The ideal location must be able to provide 45 MW of electrical demand, and to curtail costs, renewable energy should be incorporated within the utility's electric generation mix. Heavy access to the city's water system is imperative: The building requires 550,000 gallons of water per day and outputs 350,000 gallons of wastewater per day. The location must be in an EPA air quality attainment district or zone across all criteria pollutants. The location should also be near industrial railways and interstate highways. Due to the nature of the manufacturing processing requirements, the ideal location should be situated to minimize the risk of natural disasters. 350 new jobs will be created, so access to local training programs and community college collaborations is a priority. Because the brand is internationally recognized, the new location should adhere to DEI-friendly legislation and labor regulations. The location's economic incentives should be optimized to offset real and personal property tax costs. Rank 20 locations according to these desirables and justify the ranking of each location that is listed.
3. List 20 locations that you rejected and explain why.
4. Can you give me the raw data you used to rationalize <<INSERT CITY>> as the top choice?



Based on the inputs provided, the software-generated results are as follows when entered as a piecemeal query. Again, despite the academic consultant's multiple attempts, location results were unobtainable with the piecemeal approach using ChatGPT 3.5, and instead produced a message that recommended using a site selection consultant to perform the search.

- Seven of the 11 metros on the consultant-generated list did not appear on either the ChatGPT 4.0 or the Bard-generated lists.
- Both Bard and ChatGPT 4.0 each identified two locations that also appeared on the consultant-generated list, but they were not the same two in both chatbot lists.
- **None of the top three shortlisted locations identified by the consultant appeared in either the Bard or the ChatGPT 4.0 lists.**

Table 3. Industrial Shortlist Results Generated by the AI Platforms and the Piecemeal Approach versus the Consultant-Generated Shortlist

RANK	ChatGPT 3.5 (piecemeal)	ChatGPT 4.0 (piecemeal)	Bard (piecemeal)	Human Generated
1	N/A	Charlotte	Savannah	Bowling Green
2		Austin	Raleigh	Clarksville
3		Columbus	Austin	Louisville
4		Nashville	Denver	Fort Knox
5		Raleigh-Durham	Nashville	Cincinnati
6		Des Moines	Salt Lake City	Nashville
7		Atlanta	San Antonio	San Antonio
8		Phoenix	Indianapolis	Dayton
9		Salt Lake City	Charlotte	Springfield
10		Albuquerque	Pittsburgh	Knoxville
11		Richmond	Cincinnati	Indianapolis

Similarly weak results were seen using a narrative approach.

- Only two locations (Indianapolis and Columbus) on the consultant-generated list were also generated by the AI platforms.
- Each chatbot only identified one location that also appeared on the consultant-generated list.
- None of the top three locations on the consultant list appeared on any of the AI-generated lists.

Table 4. Industrial Shortlist Results Generated by the AI Platforms and the Narrative Approach versus the Consultant-Generated Shortlist

RANK	ChatGPT 3.5 (narrative)	ChatGPT 4.0 (narrative)	Bard (narrative)	Human Generated
1	Houston	San Antonio	San Antonio	Bowling Green
2	Atlanta	Columbus	Charlotte	Clarksville
3	Chicago	Reno	Phoenix	Louisville
4	Phoenix	Allentown	Austin	Fort Knox
5	Dallas	Greenville	Dallas	Cincinnati
6	Denver	Lincoln	Houston	Nashville
7	Charlotte	Des Moines	Indianapolis	Indianapolis
8	Indianapolis	Grand Rapids	Atlanta	Dayton
9	Las Vegas	Tulsa	Columbus	Springfield
10	Kansas City	Bakersfield	Memphis	Knoxville
11	Seattle	Albuquerque	Raleigh	Indianapolis

Unlike the office project, in the industrial project there is very little overlap between the human-generated list and the AI-generated lists. Given the complexity of industrial projects, including infrastructure requirements, it is much more difficult for the software to understand the nuances of these project requirements and respond accordingly.

Overall Performance

This study identified several key areas of the site selection process that presented challenges for AI models and reinforces that expert human input and oversight is critical. Chief among these challenges is the inconsistency of the ranked results among the platforms, which appeared to vary wildly, particularly for lower-ranked cities (#11 to #20), even when given the same set of prompts. Moreover, the final results elicited from the chatbots were not repeatable most of the time, which calls into question the “black box” analytical process that underpins the AI recommendations. On the other hand, the bots were highly efficient when providing users with an initial first pass of candidate cities as plausible options, even if the rationale behind the provided lists is unknown.

Specific areas where the software performed well and where they fell short are outlined below.

Inconsistency of Results: This is perhaps the most significant challenge to using generative AI for the site selection work. Case studies were run multiple times against the same input, and each time, it produced different rankings with justifications. The inconsistency of results appears to be significantly more affected by the input types (narrative form vs. piecemeal form), even though the prompts are ultimately equivalent. This suggests that the behavior of the bots can be quite inconsistent and is highly dependent on the whims and impulses of user interaction.

Query Size Matters: Entering a query in long monolithic narrative input yields more pertinent results. In one case, it was the only way to elicit any response. It is unclear why this is the case but we suspect that long queries may supply the bots with more immediate context and background for analysis and is more effective than chatting back and forth interactively.

Loss of Short-Term Memory: When entering prompts piecemeal (stepwise, one query at a time), it was not always clear that the bots were remembering all the earlier prompts (i.e., the bots seemed to apply more weight to the more recent prompts, and sometimes observe short-term amnesia). This is particularly true in ChatGPT 3.5 and Bard. The chatbots also tended to lose track of the original task. For instance, Bard listed the “San Francisco Bay Area” as one of the top 20 metros for the software company’s relocation out of San Francisco.

ChatGPT 4.0 is the Best “Analyst”: Based on this study, ChatGPT 4.0, the paid version of ChatGPT, is far more impressive than the free ChatGPT 3.5. While it is a bit slower at times, it seemed to provide consistently more nuanced results, and it also appeared to retain the memory of the current conversational context better. On top of a generally higher quality of results, the paid version also allows different input types.

Importance of Google: The rationales given for each of the recommended metros provided decent abstract summaries. However, none of the tools were able to provide live, on-the-ground information, such as being able to curate a list of “three move-in ready offices” for each city or up-to-date utility costs. The closest to being able to access live information is Google Bard, which exploits its wide web of tools that are real-time and have been tested by millions of users per day. We could not verify that building recommendations were, in fact, vacant, but there was an impressive attempt. Currently, only Bard is capable of presenting any live information. At the time of writing, ChatGPT 3.5 and 4.0 had stale-dated information by 7-8 months.

Human Oversight is Critical: AI models train on any data they can get their hands on, and to their own admission, they are incapable of fact-checking data before offering query results. That means they may generate imprecise, inaccurate or completely false results without knowing. If the user is not an industry expert, then they would not be able to detect false or misleading information. Even with incremental improvements to the bots, it is suspected that there will always be a need for human oversight and quality-checking.

Vulnerability to Bias and Manipulation: Tied to the previously identified limitation, the bots are likely susceptible to manipulation (or “optimization”). One could envision a scenario in which biased articles and web pages that put a certain city, metro, state, or country in either a positive or negative light are injected into the web to either help or hurt attempts to attract investment. It is unclear to what extent that campaign would bias the AI tools’ general results towards putting that location in a more positive light. Such “optimization” campaigns already exist for search engines (SEO) and social media (networks of paid follower-bots). It seems likely that tools like chatbots would also be vulnerable to the same manipulation tactics.

Lack of Geographic Insight: While AI systems are large language models with potential access to massive databases of location information, it was not apparent that the bots had an interpretative understanding of the geography of place competitiveness, which is a critical component of the site selection process. The high-level place recommendations did not seem to include much or any understanding of how places differentiate between metros, counties, and cities. Examples of specific data that site selection consultants need to assess a location include defining how urban areas sprawl across state lines; identifying utility service territories; recognizing the effects of transportation on workforce access; and calibrating dozens of cost factors that vary across overlapping jurisdictions.


As noted, in addition to variations in the final results, there were also variations in how each platform performed. The free version of ChatGPT (3.5) performs the worst overall. Its attention to detail and its knowledge base seemed lacking compared to the human results, to ChatGPT 4.0, and even to Google Bard at times. This was especially true for the more complex manufacturing plant expansion site. In particular, when entering the prompts piecemeal (sentence-by-sentence), the results yielded nothing at all.



Conclusion

Site selection projects require complex and nuanced analyses. Given the high-stake outcomes, new methodologies must be carefully evaluated. Members of The Site Selectors Guild recognize the importance of understanding how AI tools can be successfully leveraged without sacrificing the quality and accuracy of site location searches. Based on the Guild's survey results, there are several consultants introducing AI to enhance efficiencies in administrative and processing tasks. While there is evidence that AI can be a tool in the process, there is also considerable evidence that, in its current state, human oversight and insights will continue to be imperative to achieve a successful outcome. Guild members have no doubt that AI will be increasingly used in the site selection industry in the future, but given the results of the project, the results are not consistent with the outcomes that are generated by consultants with a strong history of institutional and industry knowledge.

The bots certainly have useful features, but retaining humans in the loop remains imperative for the initial setting of strategy options, identifying and weighing critical location factors, defining risk considerations, structuring the process of evaluating both quantitative and qualitative data, fact-checking data results, and guarding against AI bias and misdirection. This speaks only to some of the initial “desktop” or computer-based analytical work that must be done before proceeding with the essential field research, in-market tours, interviews, meetings, and negotiations that are required as the site selection process continues. Even desktop analysis requires the oversight of skilled professionals, experts who understand the parameters and specifications that drive project success, people who understand the geography of business competitiveness, and individuals with experience in the complexities of business location strategy decisions.





SITE
SELECTORS  RS
GUILD

www.siteselectorsguild.com