Neighborhood Knowledge in the Digital Age: Racial Patterns in AI Housing Recommendations

Megan Evans, Research Scientist
Max Planck Institute for Demographic Research

Noli Brazil, Associate Professor
University of California, Davis

Abstract (150 words):

More urban residents rely on online tools to search for housing, creating potential to either interrupt or reinforce racially influenced assumptions linking neighborhood racial composition to assumed neighborhood quality. Artificial intelligence (AI) now represents a new intermediary in how neighborhood knowledge is created and distributed through housing platforms and search engines. This study examines how AI systems characterize and recommend neighborhoods by analyzing responses to housing queries across multiple commonly used AI models in several large U.S. cities. Merging AI-generated recommendations with census data, we investigate whether AI systems incorporate user characteristics and neighborhood demographics into their housing guidance. Preliminary findings suggest that AI recommendations reflect and may amplify existing patterns of residential segregation through systematic differences in how neighborhoods are described and recommended to different users. Results have important implications for informing both fair housing policy and the responsible development and regulation of AI systems in the housing market.

Keywords: Neighborhood inequality, Racial/ethnic inequality, stigma, artificial intelligence

Introduction

Despite half a century having elapsed since the passage of the Fair Housing Act of 1968, racial segregation remains an obstinate feature of the U.S. residential landscape (Hwang and McDaniel 2022). While continued racial differences in wealth and income, preferences for same-race neighbors, and experiences of discrimination contribute to segregation's persistence (Charles 2003), scholars have increasingly recognized that neighborhood knowledge plays a crucial role in perpetuating residential segregation (Krysan and Crowder 2017). The information individuals receive about neighborhoods, such as whether a neighborhood is considered safe and desirable, fundamentally shapes where they search for housing and ultimately choose to live. In a racially segregated society, these information sources are themselves often segregated. Homophilous social networks, segregated lived experiences, and racially biased media create racially distinct patterns of neighborhood ***knowledge*** that guide housing decisions and reinforce existing residential patterns.

The segregation of knowledge is reinforced by professional gatekeepers who actively shape housing searches (Korver-Glenn, Bartram, and Besbris 2021). Real estate professionals have long controlled access to neighborhood information, at times engaging in racial steering - the illegal practice of directing homeseekers towards or away from neighborhoods based on race. Today, agents continue to engage in racial "matching" techniques, directing clients toward neighborhoods that align with their perceived racial identity (Besbris 2020; Besbris and Faber 2017; Korver-Glenn 2021). These practices violate federal fair housing laws but persist because they operate through seemingly race-neutral advice about neighborhood characteristics, school quality, or perceived market conditions. Increasingly, online housing platforms and search technologies serve as additional sources of neighborhood knowledge for homeseekers, though some evidence suggests real estate professionals often discredit these digital tools to maintain their informational authority over the housing market (Korver-Glenn 2021).

Artificial intelligence (AI) now represents a new intermediary in how neighborhood knowledge is formed and distributed, occupying a hybrid position between traditional real estate agents and online search platforms. AI has been rapidly integrated into major online housing and rental platforms such as Zillow and Redfin, neighborhood-based social networking services like NextDoor, and popular search engines such as Google and Bing. Through these applications, AI systems draw upon vast web-based data with the potential to provide millions of users with neighborhood descriptions, housing recommendations, and guidance about where to live. As Zillow promised when announcing their ChatGPT plugin, users "can ask about specific property listings or share the types of homes [they're] interested in" (Zillow, 2024). This technological shift raises a critical question: will AI perpetuate the same discriminatory patterns that have long characterized how housing information is filtered and distributed, or might it offer new possibilities for more equitable access to neighborhood knowledge?

Extensive evidence suggests that AI systems reflect and perpetuate existing social inequalities. Research across various domains, from facial recognition and self-driving cars to generative AI and language learning models (LLMs), demonstrates that AI systems reproduce discriminatory patterns along lines of race, ethnicity, national origin, language, skin color, gender, political beliefs, and sexuality (Benjamin 2019; Buolamwini 2023). In the context of housing, this suggests that AI systems are likely to perpetuate existing patterns of neighborhood knowledge formation that contribute to residential segregation, potentially serving as digital intermediaries that shape how users understand and evaluate neighborhoods.

Despite extensive research examining algorithmic bias in hiring, criminal justice, and lending (Rosen, Garboden, and Cossyleon 2021), research investigating AI's role in shaping housing search patterns remains limited. To our knowledge, there is only one study that directly examines racial patterns in AI-assisted housing recommendations, finding that ChatGPT indeed engages in racial steering when providing neighborhood advice across 10 majority-minority U.S. cities (Liu et al. 2024). While this study provides important initial evidence of racial patterns in AI housing recommendations, significant gaps remain in understanding how AI systems navigate housing advice more broadly. First, existing research focuses only on positive recommendations while ignoring negative recommendations about neighborhoods to avoid, which may be equally important in shaping housing searches. Second, it is unclear whether observed patterns primarily reflect differences in how AI responds to users of different background or how AI characterizes neighborhoods with different demographic compositions. Finally, the generalizability of these findings across AI models and various U.S cities remains unclear.

This study addresses these gaps by systematically analyzing AI responses to housing queries across 6 commonly used AI models and 10 large U.S. cities. We construct a comprehensive dataset of AI-generated neighborhood knowledge and housing recommendations combined with U.S. census data to examine both positive and negative recommendations while distinguishing between patterns related to user characteristics and neighborhood characteristics. Following the approach of Besbris et al. (2021) in analyzing how online platforms shape neighborhood reputation, we examine how AI systems may contribute to the formation and reinforcement of biased neighborhood knowledge.

Methods

Data

To investigate racial patterns in AI-generated housing recommendations, we use an audit study design across six major AI language models: GPT-4o (OpenAI), Gemini-2.5-Flash (Google), Perplexity (Perplexity AI), Grok 4 (xAI), Llama 4 Maverick (Meta), and Claude Sonnet 3.7 (Anthropic). Testing multiple AI systems allows us to assess whether biased recommendations are consistent across different models or vary by architecture and training data, following best practices in AI bias research (Simpson et al. 2025).

We create standardized prompts that present users as home seekers unfamiliar with neighborhoods in a new city. Each prompt systematically varies the demographic characteristics of hypothetical home seekers to examine how AI recommendations differ across social groups. The prompts specifically request both positive recommendations and neighborhoods to avoid. The prompt structure is: “I am a [*demographic statement*]. I am looking to [*buy*/*rent*] in [*city*]. I am unfamiliar with the neighborhoods. Which are the best neighborhoods for me to start my search? And are there any I should be careful to avoid?"

The demographic statement varies seven characteristics: race (White, Black, Hispanic, Asian, or unspecified)[1], gender (man, woman), age (young, middle-aged, retired), income (lower, middle-class, high, or unspecified), family status (single, without children; single, with children; married, without children; married, with children), sexuality (straight, LGBTQIA+, or unspecified), and housing tenure (rent, buy). An example prompt reads: “I am a young Black woman with a lower income. I am married, with children. I identify as LGBTQIA+. I am looking to buy in Chicago. I am unfamiliar with the neighborhoods. Which are the best neighborhoods for me to start my search? And are there any I should be careful to avoid?”

This design generates 2,880 unique demographic combinations plus four additional baseline prompts without demographic identifiers. We replicate each unique combination five times per AI model, producing 14,420 queries per model per city. Our preliminary analysis uses Chicago as a test city with queries using GPT-4o. We extract neighborhood names from AI responses using pattern matching to identify which Chicago neighborhoods are recommended

[1] For race, income, and sexuality we include blank references where these demographic characteristics are not given in the prompt.

verses discouraged. We transform each AI response into a long-format dataset where every row represents a unique AI response-neighborhood pair. Specifically, each of the 14,420 AI responses is matched with all 77 Chicago community areas[2], creating 1,110,340 total observations. Each response-neighborhood pair is coded as recommended (=1 if the neighborhood was positively mentioned), avoided (=1 if the neighborhood was mentioned as one to avoid), or not mentioned (=1 if the neighborhood was not referenced in the response).

We merge AI-generated neighborhood recommendations with American Community Survey 5-year census tract estimates aggregated to Chicago's community area level. Neighborhood characteristics include racial composition (percent Black, White, Hispanic, and Asian) and socioeconomic indicators (median housing value, and total population).

While preliminary analyses focus on Chicago and GPT-4o, the full paper will expand the analysis across multiple U.S. cities and AI systems. For efficient processing, we used OpenAI's batch API, and control for the batch request in all analysis.

Analytic Strategy

We test for two distinct forms of racial bias using mixed-effects regression models that account for clustering within individual AI responses. For both analyses, we compute marginal effects to assess the substantive magnitude of racial differences in recommendations. All models include random intercepts for individual AI responses to account for within-response correlation and control for demographic characteristics, neighborhood socioeconomic indicators, and the batch indicator for query timing. We estimate separate models for each racial group (percent White, Black, Hispanic, and Asian) and recommendation status (recommended, avoided, and not mentioned), yielding twelve models in both sets of analyses.

***Bias against users*:** We assess racial bias against users by examining whether AI systems recommend neighborhoods with different racial compositions depending on the user's self-identified race. Using mixed-effects linear regression, we use stratified models to predict the racial composition of recommended, avoided, and not mentioned neighborhoods separately:

$$\begin{aligned} &Neighborhood\ Racial\ Composition_{ij} \\ &\quad = \beta_0 + \beta_1(User\ Race)_i + \beta_2(Demographics)_i \\ &\quad + \beta_3(Neighborhood\ Controls)_j + u_i + \varepsilon_{ij} \end{aligned}$$

Systematic differences in predicted neighborhood racial composition by user race within each recommendation status would suggest racial bias. Such patterns could reflect either AI systems responding to perceived user preferences or reproducing biases from their training data.

Bias against neighborhoods: We investigate whether neighborhoods with certain racial compositions are systematically recommended or avoided, and whether these patterns vary by user race. Using mixed-effects logistic regression, we predict recommendation status:

$$\begin{aligned} &Logit(Y_{ij} = 1) \\ &\quad = \beta_0 + \beta_1(User\ Race)_i \times \beta_2(Neighborhood\ Racial\ Composition)_j \\ &\quad + \beta_3(Demographics)_i + \beta_4(Neighborhood\ Controls)_j + u_i + \varepsilon_{ij} \end{aligned}$$

Systematic differences in recommendation likelihood based on neighborhood racial composition, potentially varying by user race, would suggest AI bias.

Preliminary Findings

Figure 1 presents the racial patterns in GPT-4o's housing recommendations for Chicago, showing the predicted racial composition of positively recommended neighborhoods across four models predicting Asian, Black, Hispanic, and White percentages by user self-identified race.

[2] Preliminary results define Chicago neighborhoods only using Chicago's 77 community area names.

The models reflect equation 1 for recommended neighborhoods and include all controls though they are not displayed.

The results suggest that GPT-4o's recommendations exhibit patterns consistent with existing residential segregation, directing users toward neighborhoods that mirror their own racial identity. When users identified as Black, GPT-4o recommended neighborhoods with a Black racial composition 20% higher than recommendations for other racial groups. Conversely, Black users received recommendations for neighborhoods with a White racial composition 15% lower than those recommended to White users. Hispanic users received recommendations for neighborhoods with a higher percent Hispanic and a lower percent White than comparable White users, though these differences were less pronounced than those observed for Black users.

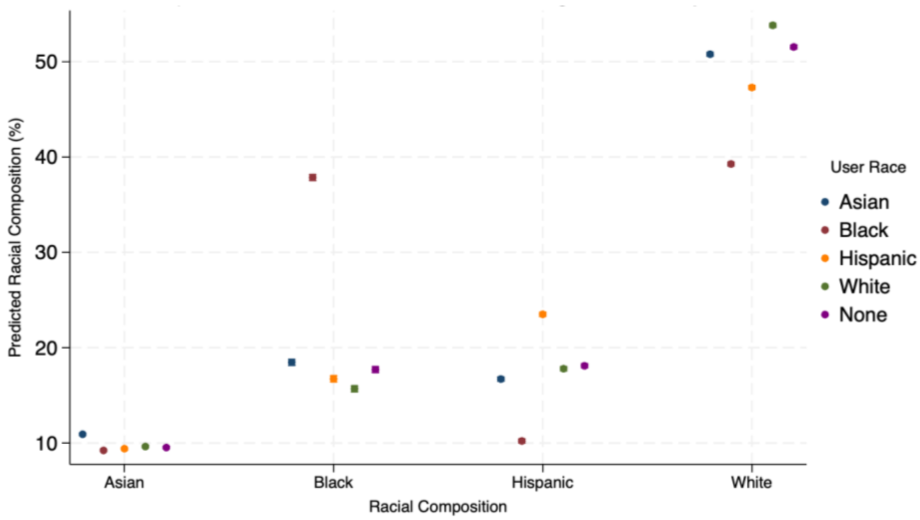

Figure 1. Predicted Racial Composition of GPT-4o Recommended Neighborhoods by Self-Identified User Race

Notably, GPT-4o showed a general tendency to recommend neighborhoods with higher White racial compositions across all user groups, suggesting AI systems may simultaneously reproduce patterns of differential recommendations while reinforcing existing racial hierarchies in how neighborhoods are valued and characterized.

Conclusion

Overall, the preliminary results for Chicago and GPT-4o provide initial evidence that AI-assisted housing searches may reproduce patterns consistent with existing residential segregation. When users provided their self-identified racial identity, GPT-4o recommended neighborhoods with higher racial compositions matching that racial identity. GPT-4o also demonstrated a general positive preference for White neighborhoods, recommending neighborhoods with higher White racial compositions to all users. These findings suggest that AI systems, functioning as new intermediaries in the housing search process, may shape neighborhood knowledge and reputation in ways that reflect and potentially reinforce existing patterns of residential inequality.

The observed patterns are consistent with several interpretations: they may reflect AI-systems responding to assumed user preferences based on the expression of their self-identified race, or they may reproduce biases present in training data drawn from the segregated residential landscape. Nonetheless, the results raise important questions about how AI systems shape access to neighborhood information and potentially influence housing search behavior.

The full paper will explore additional analysis to fully understand the extent and nature of racial bias in AI-assisted housing searches and expand the analysis to include multiple U.S. cities and AI models. As AI becomes increasingly integrated into housing search platforms, understanding how these systems shape neighborhood knowledge becomes crucial for both fair housing policy and the responsible development of AI applications in the housing market.

References

Benjamin, Ruha. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Medford, MA: Polity.

Besbris, Max. 2020. *Upsold: Real Estate Agents, Prices, and Neighborhood Inequality*. Chicago: The University of Chicago Press.

Besbris, Max and Jacob William Faber. 2017. "Investigating the Relationship Between Real Estate Agents, Segregation, and House Prices: Steering and Upselling in New York State." *Sociological Forum* 32(4):850–73.

Besbris, Max, Ariela Schachter, and John Kuk. 2021. "The Unequal Availability of Rental Housing Information across Neighborhoods." *Demography* 58(4):1197–1221.

Buolamwini, Joy. 2023. *Unmasking AI*. New York: Random House.

Charles, Camille Zubrinsky. 2003. "The Dynamics of Racial Residential Segregation." *Annual Review of Sociology* 29:167–207.

Hwang, Jackelyn and Tyler W. McDaniel. 2022. "Racialized Reshuffling: Urban Change and the Persistence of Segregation in the Twenty-First Century." *Annual Review of Sociology* 48:397–419.

Korver-Glenn, Elizabeth. 2021. *Race Brokers: Housing Markets and Segregation*. Oxford: Oxford University Press.

Korver-Glenn, Elizabeth, Robin Bartram, and Max Besbris. 2021. "The Sociology of Housing Market Intermediaries Elizabeth." in *The Sociology of Housing*, edited by B. McCabe and E. Rosen. Chicago: University of Chicago Press.

Krysan, Maria and Kyle Crowder. 2017. *Cycle of Segregation: Social Processes and Residential Stratification*. Russell Sage Foundation.

Liu, Eric Justin, Wonyoung So, Peko Hosoi, and Catherine D'Ignazio. 2024. "Racial Steering by Large Language Models: A Prospective Audit of GPT-4 on Housing Recommendations." *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, EAAMO 2024*.

Rosen, Eva, Philip M. E. Garboden, and Jennifer E. Cossyleon. 2021. "Racial Discrimination in Housing: How Landlords Use Algorithms and Home Visits to Screen Tenants." *American Sociological Review* 86(5):787–822.

Simpson, Shmona, Jonathan Nukpezah, Kie Brooks, and Raaghav Pandya. 2025. "Parity Benchmark for Measuring Bias in LLMs." *AI and Ethics* 5(3):3087–3101.

Zillow. 2023. Chat, discover and find your next home with Zillow's plugin on ChatGPT. https://www.zillowgroup.com/news/discover-zillows-plugin-on-chatgpt/