

## Toward AI Matching Policies in Homeless Services: A Qualitative Study with Policymakers

Caroline M. Johnston<sup>1</sup>, Olga Koumoundouros<sup>1</sup>, Angel Hsing-Chi Hwang<sup>1,2</sup>, Laura Onasch-Vera<sup>1</sup>,  
Eric Rice<sup>1</sup>, Phebe Vayanos<sup>1,3,4</sup>

<sup>1</sup>CAIS Center for AI in Society, University of Southern California

<sup>2</sup>Annenberg School for Communication, University of Southern California

<sup>3</sup>Department of Industrial and Systems Engineering, University of Southern California

<sup>4</sup>Department of Computer Science, University of Southern California

carolinemjohnston39@gmail.com, {koumound, hsingchi, onaschve, ericr, vayanou}@usc.edu

### Abstract

Artificial intelligence researchers have proposed various data-driven algorithms to improve the processes that match individuals experiencing homelessness to scarce housing resources. It remains unclear whether and how these algorithms are received or adopted by practitioners and what their corresponding consequences are. Through semi-structured interviews with 13 policymakers in homeless services in Los Angeles, we investigate whether such change-makers are open to the idea of integrating AI into the housing resource matching process, identifying where they see potential gains and drawbacks from such a system in issues of efficiency, fairness, and transparency. Our qualitative analysis indicates that, even when aware of various complicating factors, policymakers welcome the idea of an AI matching tool if thoughtfully designed and used in tandem with human decision-makers. Though there is no consensus as to the exact design of such an AI system, insights from policymakers raise open questions and design considerations that can be enlightening for future researchers and practitioners who aim to build responsible algorithmic systems to support decision-making in low-resource scenarios.

**Extended version** — <https://arxiv.org/abs/2508.07129>

### Introduction

As of January 2024, over 75,312 persons are experiencing unsheltered homelessness in Los Angeles (LA) (Los Angeles Homeless Services Authority 2024a). However, only 22,266 permanent housing units exist in the housing system, of which most are already occupied (Los Angeles Homeless Services Authority 2024b). Los Angeles uses the Vulnerability Index – Service Prioritization Decision Assistance Tool (VI-SPDAT), a 35-question self-reported survey about an individual’s housing history; risks; socialization and daily functioning; and wellness, as part of prioritizing individuals for these scarce housing resources (Orgcode 2014). The VI-SPDAT measures an individual’s vulnerability, where, in general, those who are more vulnerable (i.e., at higher risk of ‘adverse outcomes’) are prioritized by human decision-makers who match individuals to resources and services.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Though its use was well-intentioned, the VI-SPDAT has received great scrutiny in recent years due to evidence of racial bias within its administration and scoring results (Los Angeles Homeless Services Authority 2018). Additionally, the tool was never designed to match individuals to resources (De Jong 2024), resulting in misalignment between the system’s goals and its use.

Concurrently, with the recent explosion of generative artificial intelligence (AI) tools such as ChatGPT,<sup>1</sup> the general public is recognizing the power of AI in creating content, making predictions, and aiding in decision-making (Fried 2024). For example, in the homelessness field, AI researchers have created data-driven methods that prioritize or allocate scarce resources to individuals experiencing homelessness, potentially accounting for fairness (e.g., less racial bias) in the process (Azizi et al. 2018; Toros and Fleming 2018; Kube, Das, and Fowler 2019; Rahmattalabi et al. 2022; Kube, Das, and Fowler 2023; Kumar and Yeoh 2023; Tang et al. 2023). However, the public has also become more aware of the harm caused by AI. Many negative perceptions of AI, including its ability to reinforce bias and its lack of transparency, persist within social services, creating a reluctance to adopt AI tools (Denton 2019; Eubanks 2019; Reamer 2023). It remains to be seen whether AI tools that attempt to address potential fairness and transparency risks would be accepted within the homeless services community.

In this work, we aim to answer the following questions: *Are homeless services policymakers in Los Angeles open to the use of AI-driven decision-support tools for matching individuals experiencing homelessness to housing resources? What are the perceived benefits and harms of pursuing this idea? What difficulties would lie ahead?* To this end, we conduct semi-structured qualitative interviews with 13 policymakers who directly influence the prioritization policies in the housing allocation system of LA to understand their perceptions of utilizing AI tools within the matching system.

Using the AI allocation policies of Tang et al. (2023) as a case study, we present potential efficiency, fairness, and transparency features of such a system and receive feedback on these topics. We focus on this particular AI system because of its use of causal inference techniques to optimally assign individuals to resources to maximize the

<sup>1</sup><https://chatgpt.com/>

number of individuals that successfully exit homelessness, subject to resource and fairness constraints. This aims to achieve a balance between efficiency and fairness in this highly constrained environment. Though it is not the only AI method in the homelessness field that uses causal inference (Kube, Das, and Fowler 2019; Rahmattalabi et al. 2022; Kube, Das, and Fowler 2023; Kumar and Yeoh 2023), it offers a transparent and responsive allocation policy, optimally assigning individuals whenever an individual or new resource enters the system, while others solve at a less granular and potentially delayed level, i.e., assigning resources on a weekly (Kube, Das, and Fowler 2019, 2023) or monthly (Kumar and Yeoh 2023) basis.

As interest in research that informs AI governance grows, this work aims to guide how technologists can collaborate with policymakers to design systems that address real-world policy needs in the homeless services domain. Specifically, we make the following contributions:

1. We reveal Los Angeles policymakers' openness to using AI tools to assist human decision-makers in the matching process and identify additional AI applications they view as beneficial for homeless services beyond matching.
2. We identify potential gaps between what AI matching tools deliver and policymakers' desired outcomes for fairness, efficiency, and transparency, based on a case study of a specific AI tool and interviews with policymakers in homeless services.
3. We additionally identify potential implementation roadblocks for such systems.

The current study served as one of the very few that directly engaged with policymakers to extract actionable insights, complementing most prior work that studied perspectives of frontline workers and individuals experiencing homelessness (Kuo et al. 2023; Tang et al. 2024; Tracey and Garcia 2024a,b; Moon et al. 2025). We also present recommendations and insights for AI researchers interested in bringing AI tools into these spaces.

## Literature Review

Various works involving algorithmic decision-making (i.e., AI) use qualitative methods to elicit feedback and insights on the design of such tools from relevant stakeholders. In public services, these studies range from child welfare (Chouldechova et al. 2018; Saxena et al. 2021; Stapleton et al. 2022; Saxena and Guha 2024) to job placement (Møller, Shklovski, and Hildebrandt 2020; Flügge, Hildebrandt, and Møller 2021), to criminal justice (Brayne and Christin 2021), to refugee replacement (Ahani et al. 2021). For example, Lee et al. (2019) develops a matching algorithm with a nonprofit that provides on-demand donation transportation through the use of volunteers, exploring the equity-efficiency tradeoff of such an algorithm. Both Kuo et al. (2023) and Stapleton et al. (2022) use co-design activities to incorporate relevant stakeholders preferences and feedback into algorithmic processes in homeless services and child welfare, respectively.

In recent years, AI researchers have studied how to integrate AI tools into the delivery of homeless services. Toros

and Flaming (2018) develop a predictive algorithm to target and prioritize unhoused individuals who are predicted to have high future costs for public services. The authors argue that by prioritizing these high cost individuals for the most supportive services, then the incurred savings can be allocated to low cost individuals. Both Kube, Das, and Fowler (2019) and Tang et al. (2023) use causal inference techniques on observational data to design a resource allocation system for individuals experiencing homelessness. Kube, Das, and Fowler (2019) focuses on allocating resources to reduce the number of families who repeatedly enter homelessness while Tang et al. (2023) assigns clients to queues for resources to maximize the number of individuals that exit homelessness subject to potential fairness constraints. Taking the perspective of how to design prioritization policies subject to stakeholder preferences, Vayanos et al. (2020) designs a robust online preference elicitation algorithm and applies it to the setting of learning policymakers' preferences for potential trade-offs in efficiency and fairness metrics of housing allocation policies.

Some have specifically studied algorithmic fairness with applications to homeless services. Mashiat et al. (2022) examines fairness trade-offs in scarce resource allocation applied to administrative records from homeless services while Akpinar, Lipton, and Chouldechova (2024) studies the fairness of public sector decision-making algorithms using under-reported administrative data. Watson-Daniels et al. (2023) studies fairness inspired by scoring rules used to prioritize individuals experiencing homelessness for housing resources in Allegheny County (See Showkat et al. (2023); Moon and Guha (2024) for more complete surveys of AI in homeless services applications.)

The works listed above study methodological approaches of how to "best" or "fairly" use data to prioritize or allocate resources to individuals experiencing homelessness and do not interface with those using or affected by the proposed AI tools. Thus, Human-Computer Interaction (HCI) researchers have begun to engage directly with the community in how its members interact with data and AI tools in homeless services (Karusala et al. 2019; Kuo et al. 2023; Tang et al. 2024; Tracey and Garcia 2024a,b; Moon et al. 2025). For example, Kuo et al. (2023) engages with service providers and unhoused individuals about an already deployed data-driven prioritization tool that uses government administrative records instead of the typical, self-reported information from clients about their current housing situation. Tang et al. (2024) conducts workshops with service providers and unhoused individuals about "AI Failure Cards," a way for the community to understand and provide feedback on the failures of a deployed AI housing allocation algorithm.

Our work differs from these listed above in that we engage with policymakers, not frontline workers, in the discussion of the design of an AI allocation tool in the homeless services field. Because policymakers are more accustomed to "systems level" thinking, they are uniquely positioned to comment on topics such as the challenges of implementing AI across their organizations and the homeless services ecosystem. By using an AI resource allocation model case study to ground our discussion, policymakers are able to pro-

vide feedback that can be addressed or pursued head-on as opposed to post-deployment, if such a system were built.

## The Los Angeles Housing System

We now provide the necessary background of the LA housing resource allocation system.

### Vulnerability Assessment

As a requirement from the United States Housing and Urban Development (HUD), the Los Angeles Housing Services Authority (LAHSA) has developed policy to prioritize and match individuals experiencing homelessness (clients) to available resources by their level of vulnerability (U.S. Department of Housing and Urban Development 2017). To assess vulnerability, many housing service providers in LA will ask individuals to take the VI-SPDAT. From the client's answers to the VI-SPDAT, they receive points that gauge their risk in a score of 0 to 17, where 0 is the lowest level of vulnerability and 17 is the highest level of vulnerability. In general, the survey is administered to a client at locations that help to connect them to housing services, such as shelters or through street outreach.

Historically, clients' VI-SPDAT scores were used to create an ordered priority queue to receive scarce housing resources. For example, a score between 8-17 would prioritize an individual by score for Permanent Supportive Housing (PSH) (e.g., affordable housing, case management services). In recent years, however, the community and researchers in homeless services have identified numerous issues in using the VI-SPDAT this way. This includes a racial bias in which White clients tended to score higher than non-White clients and findings that the VI-SPDAT does not accurately capture an individual's true level of vulnerability (Los Angeles Homeless Services Authority 2024c). Because of these complicating factors and a desire to prioritize individuals that are truly the most vulnerable, in more recent years, LAHSA's prioritization policy uses the VI-SPDAT score as an eligibility factor instead of an ordered priority and a revised version of the VI-SPDAT that more accurately assesses vulnerability (Rice et al. 2023).

A client's VI-SPDAT answers and score, with other information such as their demographics and location, are recorded within the Home Management Information System (HMIS), an online database for collecting and accessing data on homelessness and housing resources. The HMIS database also records any resource that an individual is matched or connected with and attempts to track client outcomes.

### Matching Clients to Housing Resources

LA is divided into eight geographic regions called Service Planning Areas (SPAs). Each SPA is responsible for servicing its clients with its designated housing resources. Historically, each SPA had a small number of *matchers* who matched clients to available housing resources on an individual basis, taking into account their VI-SPDAT score, information from their case worker, and eligibility requirements of the housing resource. For example, some PSH units are

only available for veterans or those with an HIV/AIDS diagnosis, as set by the unit's particular funding source.

As of July 2024, LAHSA has assumed a centralized responsibility for matching clients to PSH instead of the SPA matchers to streamline and increase visibility of the process (Los Angeles Homeless Services Authority 2024d). For the purposes of this work, considering an AI matching tool is still relevant and can most likely adapt to such a change.

### Known Critiques of the Current System

**Efficiency** The VI-SPDAT was not intended for matching, something its creators have repeatedly noted, though various communities across the United States have adopted it as such (De Jong 2024). Though LAHSA changed the way it is used within the matching process, as described in the previous section, the matching policy is not well-aligned in terms of improving potentially desired outcomes of the system. For example, if the goal of the policy is to get the most vulnerable individuals off the street, the VI-SPDAT is a flawed proxy to measure client vulnerability (Los Angeles Homeless Services Authority 2024c) or risk of return to homeless services (Brown et al. 2018), though there have been efforts to address its shortcomings (Rice et al. 2023).

**Fairness** Stakeholders at LAHSA and in homeless services at large are interested in promoting equity through homeless services. This is evident through LAHSA's changes to the administration and use of the VI-SPDAT in response to the evidence of racial bias in the tool. In addition to these changes, to further investigate the challenges and obstacles of marginalized demographics of clients, LAHSA assembled an Ad Hoc Committee on Women and Homelessness and the Ad Hoc Committee on Black People Experiencing Homelessness (Los Angeles Homeless Services Authority 2017, 2018). The latter committee determined that though Black clients receive PSH at equal or higher rates compared to other racial groups, Black clients return to homelessness at much higher rates. They also found that Black clients are disproportionately represented in the homeless population in LA.

**Transparency** The housing system is a vast and complicated system (see Results), especially as policies, requirements, and responsibilities have shifted in recent years. As we will discuss, it is unclear whether clients accurately understand how the matching process works and the factors that facilitate their matching. Additionally, some of these processes and information are purposely withheld from clients to protect them, manage their expectations, or prevent them from "gaming" the system.

### Methodology

We now describe the methodology of our qualitative interviews and analysis. Our study draws inspiration from various works that engage with stakeholders to gain insight into the use and perceptions of AI tools in various domains such as homelessness (Kuo et al. 2023; Tang et al. 2024), food donation transportation (Lee et al. 2019), and child welfare (Stapleton et al. 2022). Our interviews have policymak-

ers elaborate on the shortcomings of the historical and current matching policies in regards to efficiency, fairness, and transparency. We then discuss a specific AI matching policy as a case study to inform discussions of how AI could be used for matching. Finally, policymakers express their attitudes and concerns about using AI in such applications.

## Recruitment & Demographics of Interviewees

We recruit 13 participants through a contact at LAHSA. These participants consist of policymakers in homeless services who either directly advise LAHSA or are representatives from leading service providers in particular SPAs. The group includes individuals with lived homelessness experience and employees of various homeless services providers, non-profits, and government agencies in LA, forming a diverse group of stakeholders in functions, experiences, and geography. These participants directly influence LAHSA's policies and procedures, including its prioritization policies for matching individuals experiencing homelessness to housing resources.

Though small in relative size, these policymakers come from a *highly specialized and influential* group in homeless services who affect key policy decisions in LA. They have the political agency to either discourage or catalyze the adoption of AI matching tools at the system level. Thus, it is critical to gauge their openness and hesitations to such ideas.

To protect the confidentiality of participants, we present their aggregate demographic statistics in Table 1. If their affiliating organization allowed such a payment, participants received a \$100 Amazon gift card for their participation. Our Institutional Review Board (IRB) reviewed and approved the study's protocols and materials.

## Protocol

One of the researchers, who is an expert in AI and matching systems, conducted 60-75 minute semi-structured interviews with participants over Zoom. We first asked about their previous experiences with and exposure to any AI in either their personal or professional lives (e.g., ChatGPT, recommendation services such as Netflix, or voice-assistants). We then asked whether they had previously considered using AI on HMIS data to assist in making matching decisions. We additionally asked about the shortcomings of the current matching system regarding its efficiency, fairness, and transparency.

To discuss the potential benefits and drawbacks of using AI for matching more concretely, we presented the results and properties of the AI matching policies of Tang et al. (2023), trained on historical LA HMIS data, to the policymakers as a case study. The AI policies use causal inference techniques to assign clients to queues for a type of housing resource to maximize the expected number of individuals that successfully exit homelessness under the resource (or no resource) received (see the work itself for technical details). These queues include a "services only" queue (e.g., a day shelter or childcare services), a "rapid rehousing" queue (e.g., short term rental assistance), or a "PSH" queue. This assignment accounts for a limited budget or availability of each resource type and potential fairness constraints.

We describe the transparency features of the AI policy to the policymakers in the following sense: the queues act in a "first come, first served" fashion where once a client enters the system (e.g., takes the VI-SPDAT), the AI policies use the client's HMIS information to assign them to a resource queue. When an available resource arrives in the system (e.g., a new PSH unit), it is presented to the first client in that respective resource line (e.g., the PSH queue). Clients that "look the same" in the data will be placed into the same queue, receiving the same type of resource or support. Because clients are in a queue, they or their case manager can know the client's place in line. Additionally, if the arrival rate of resources can be estimated (e.g., through expected construction timelines for PSH), clients or case managers can be provided with estimated wait times.

We additionally showed policymakers the potential gains in fairness (equity) by implementing these AI matching policies compared to the "status quo," or historical matching policy. We focused on the following two metrics: 1) "allocation parity," or ensuring that the rates that individuals receive housing resources are equal across protected groups and 2) "outcome parity," or ensuring that the rates that individuals successfully exit homelessness are equal across protected groups. Particularly, we focus on race as the protected group for both metrics and the allocation rates of PSH as the resource in allocation parity. Previous work has shown that it is very unlikely for an allocation policy to achieve both allocation parity and outcome parity (Jo et al. 2023), so we present results of two AI allocation policies to the policymakers, where one enforces allocation parity and the other enforces outcome parity (see Figure 1).

We stress that our intent is not to see if the policymakers want to adopt the method of Tang et al. (2023), but to use it as a concrete example of an AI matching policy to seek feedback on the properties of such a system. We did not focus on explaining the technical details of how the AI policies create the matching (i.e., causal inference and its assumptions), but instead focus more on what the AI policies could achieve in terms of its efficiency, equity, and transparency features. Finally, we asked for the policymakers' receptiveness and hesitations on adopting and implementing AI policies within the matching system in general. Our detailed interview protocol can be found in the Appendix of the extended version.

## Qualitative Analysis

The audio transcripts from the interviews amounted to 14 hours of qualitative data. To analyze these, we used a hybrid of inductive and deductive thematic analysis and used the qualitative data analysis software ATLAS.ti.<sup>2</sup> Two of the researchers, an expert in AI and an expert in Social Work, separately coded a subset of the transcripts using open coding. They then met to discuss and resolve differences in their interpretations of the data, considering their differing expertise and perspectives across disciplines, and formed a set of initial codes and codebook (Padgett 2008). They then coded another subset of transcripts and continued this process until they reached consensus after three total rounds of collabora-

<sup>2</sup><https://web.atlasti.com/>

Demographic	Frequency (%)	Demographic	Frequency (%)
Advisory role representative	77	1-4 years professional experience	8
Lead service provider representative	23	5-9 years professional experience	23
Male	23	10-14 years professional experience	38
Female	69	15+ years professional experience	31
Non-binary	8	White	46
20-29 years of age	8	Black/African American	23
30-39 years of age	38	Hispanic/Latino	31
40-49 years of age	31	Experience as matcher or case worker	31
50+ years of age	23	Lived homelessness experience	15

Table 1: Aggregate Demographics of Participants ( $n = 13$ )

rative discussion and coding. The AI researcher then independently coded the remaining transcripts using the agreed upon codebook. Then the whole team met to give feedback on the codes and structure the themes for our thematic analysis. The finalized set of codes is listed in the Appendix of the extended version. They include those that arose directly from the interview questions, e.g., “Perceptions of AI,” or how participants felt about integrating AI into the matching system, while others emerged organically from the discussions with the participants, e.g., “Other AI Solutions,” or areas beyond matching that policymakers envisioned that AI could assist with.

### Study Site

We recognize that this study focuses on policymakers in LA and LA’s specific housing resource matching process for individuals experiencing homelessness, which may not always generalize elsewhere. For example, not every community operates with a comparable SPA system. However, we believe this study contains insights that are applicable beyond LA. For example, as of 2024, many communities still use the VI-SPDAT in the process of prioritizing and/or matching clients including Santa Clara County, California (The County of Santa Clara 2024) and parts of Canada (City of Hamilton 2024; Homeless Individuals and Families Information System 2024). The issues discussed in our results such as messy data (Kube, Das, and Fowler 2023; Kuo et al. 2023; Moon and Guha 2024; Moon et al. 2025), bias from human decision-making (Chouldechova et al. 2018; Denton 2019; Moon and Guha 2024; Moon et al. 2025), and high turnover rates in the workforce (Moses 2023; KPMG 2024) permeate across homeless services communities beyond LA. In fact, various other domains encounter these issues when implementing data-driven tools including policing and criminal courts (Brayne and Christin 2021); donation transportation services (Lee et al. 2019); refugee resettlement (Ahani et al. 2021); and job placement (Møller, Shklovski, and Hildebrandt 2020).

## Results

We stress that these participants are experts in homeless services, not in AI, and do not necessarily come from technical

backgrounds. If participants had a question about the capabilities of AI, the AI expert interviewer would answer to the best of their ability. Misconceptions of AI by the participants were not necessarily corrected.

We now present the specific findings from our analysis. In the following, we 1) assess the pre- and post-interview sentiments of participants concerning AI. We then 2) identify potential gaps between the AI matching model case study of Tang et al. (2023) and the real matching process as revealed by the policymakers. In other words, what remains to create a real AI matching system and are these requirements discussed with respect to the available data that may be used by AI. Finally, we 4) synthesize policymakers’ openness and hesitations to using AI in homeless services both in general and with respect to the AI matching case study.

### Pre- and Post-Interview Results

In Table 2, we show participants’ current level of exposure to and sentiments of AI, both in general and in terms of matching individuals experiencing homelessness to scarce housing resources. Ten participants currently use AI in their daily life. Six, or less than half, of the participants had previously thought about how AI can be used to match individuals to scarce housing resources. Ten participants had mixed sentiments about using AI for matching before the interview. All participants either had the same or a more positive perception of AI after the interview, where four participants had a more positive perception. Only one participant had a negative view in both the pre-interview and post-interview.

### Reconciling an AI Matching Model with the Complex Reality of Matching

The housing system in LA is complicated and political, as revealed below. Policymakers discussed logistical complexities within matching that are not reflected or accounted for within the AI matching case study. These include both external factors that are outside of the control of clients, case workers, and matchers; and internal complexities that must be addressed for an AI matching tool to create valid and effective matches. In fact, various policymakers envisioned a different or supplementary AI tool that can assist in some of these complex processes.

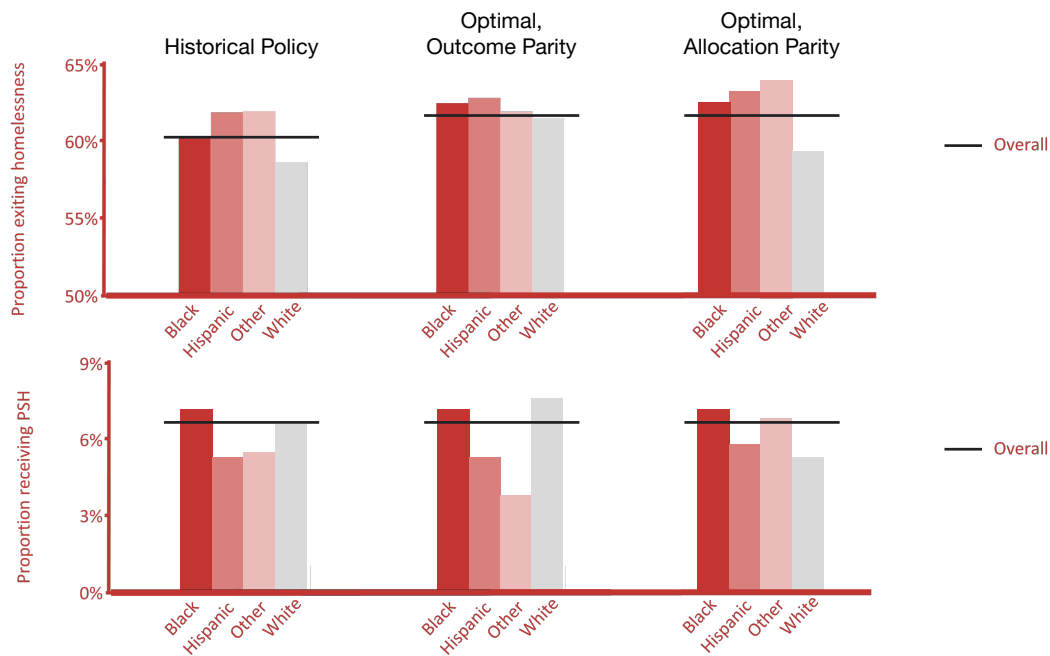


Figure 1: Fairness results of the historical allocation (“Historical Policy”) and the AI allocation policies trained to maximize the number of clients successfully exiting homelessness, showing the proportion of clients exiting homelessness and receiving PSH by race. The “Optimal, Outcome Parity” AI model ensures outcome parity across race while the “Optimal, Allocation Parity” AI model ensures allocation parity across race for PSH. The black line shows the proportion of clients exiting homelessness (resp. receiving PSH) over the entire client population. Both AI policies achieve a 2% or 300 client increase in exits from homelessness. Adapted from Tang et al. (2023).

**Requirements for matching** Policymakers mentioned many factors when housing a client, ranging from eligibility for PSH units, to whether clients have a social security card. A client also needs responsiveness from their case worker to make the match, which “*is hard because points of contact change pretty frequently*” (P6). In terms of eligibility factors for resources, these can vary greatly due to the funder of the resource, ranging from age criteria for a senior complex or for Transition Age Youth (TAY), aged 18 to 24 years old; a mental health condition; or a requirement of chronic homelessness and enrollment in a specific program such as Time-Limited Subsidies (TLS) (P10). In terms of incorporating these factors into an AI model to create valid matches, P6 doubted that all of this data exists to facilitate such a tool:

“*In theory, [the] data system has all of the different beds, the current data availability, the different eligibility criteria [...] all of that information that you need about that person. And so you can plug everything in, and it will pop out ‘This is the [housing] voucher that’s available. This is the person.’ But I don’t think we have all of that information on the system side.*” (P6)

Multiple policymakers envisioned a supplementary AI tool to help caseworkers determine what resources their client is eligible for based on the resource’s criteria and the characteristics of their client (P0, P6, P7, P11). P11 noted it

would be helpful to use AI to find the “*most amount of PSH resources*” available that a client is eligible for. P0 described a system in which AI could notify case workers of missing, required forms. With this flag, case workers could kickstart the process of gathering the appropriate information.

**No two resources nor clients are the same** Many policymakers brought up idiosyncrasies of clients or resources that are not reflected in the AI matching case study of Tang et al. (2023), which may be relevant for making matches. For example, P1 noted that in the HMIS, the exact disability and exact chronic health issues of clients are not necessarily recorded. However, these details may be relevant in regards to the level of services needed by that client:

“*Maybe [a client’s] disability is he doesn’t have one leg, but he has a prosthetic, and his health issue might be asthma, [and he has] an inhaler ... [Another client’s] disability could be schizophrenia and his chronic health issue might be COPD [Chronic Obstructive Pulmonary Disease], which is very different than asthma.*” (P1)

Currently, though this information may not be in the HMIS, case workers may be aware and make decisions accordingly.

Additionally, at an idiosyncratic level, not every client is at the same stage in their housing journey or needs the same type of support. Sometimes clients “*are not necessarily ready to be housed. Sometimes they need to go into in-*

Participant	Curr. use AI?* (Q6)	Sentiment of prev. discussions of AI* (Q6)	Prev. considered AI?† (Q7)	Sentiment of AI† pre-interview (Q8)	Sentiment of AI† post-interview (Q20)
P0	Yes	Mix	Yes	Mix+	Mix+
P1	No	Mix-	No	Positive	Positive
P2	Yes	Mix	Yes	Mix+	Mix+
P3	Yes	Mix	No	Mix+	Mix+
P4	Yes	Mix	No	Mix	Positive
P5	No	Mix	Yes	Negative	Negative
P6	Yes	Mix	Yes	Mix	Mix
P7	Yes	Mix	No	Mix+	Positive
P8	Yes	Mix	Yes	Mix	Mix+
P9	Yes	Mix	No	Mix	Mix
P10	No	Positive	No	Mix	Positive
P11	Yes	Mix	Yes	Mix	Mix+
P12	Yes	Negative	No	Positive	Positive

Table 2: Participants’ previous experience with and perception of AI ( $n = 13$ ). “+” (resp. “-”) denotes the participant indicated “leaning” towards a positive (resp. negative) sentiment when describing mixed feelings. See Appendix of the extended version for corresponding protocol question. \*AI in general. †AI for matching scarce resources to individuals experiencing homelessness.

terim housing ... [to] get a little structure in their lives. And then move on to living in a unit” (P5). P1 described how if a client has spent a long time in jail, then “a congregate setting shelter is not appropriate. It’s re-traumatizing because it reminds them of jail.” Another example is when housing a client who is a TAY: “A 19 year old is not supposed to be housed in the same place from that time until they die. They move around” (P1). Thus, some housing options are not right for every client based on their experiences. Closely related to this is incorporating client preference or choice of resources, since clients ultimately need to accept or decline any resource offered (P0, P1, P6, P8). P8 mentioned that their organization has considered how an AI tool can match resources with respect to client preferences.

The case study treats each type of resource as functionally the same, though P6 noted “PSH is not equal across the system.” For example, older units may be in “not as nice of an area that’s not as well resourced” (P6). For TLS, some subsidies are only available for a certain period of time, while others can stay with the client for an extended period (P0). P8 described differences in resources in terms of the barriers clients may encounter: “certain resources themselves are harder to use between White people and people of color.”

**External forces** Policymakers discussed various political factors that restrict the decision-making processes of the housing system. For example, the geographic boundaries of each SPA is determined by the LA County Department of Public Health. Within each SPA, there is a limited number of resources that their respective clients can access, where some SPAs do not build housing at all (P1). Thus, depending on a client’s location, it can greatly limit their resource access. From a different political aspect, many policymakers mentioned that undocumented individuals cannot receive PSH, limiting their access as well (P3, P4, P6, P7, P11).

Politicians can also have direct influence over housing

decisions, not just through legislation, but at the individual client level (P1, P3):

*“If a politician calls you and says this has to happen, we have to do it. [...] It doesn’t matter who [the client] and what their [VI-SPDAT] score is. We had to move forward with what they were wanting us to do.” (P1)*

### Data Issues: What Would AI Be Learning?

The interviews with policymakers confirmed our understanding of the fidelity and lack of trust in relevant data. Policymakers voiced many concerns about the quality of housing data due to a myriad of issues, including human error and a lack of standardized data practices across or within agencies. This lack of trust threatens the efficacy of any future AI system that would require training on historical data, whether it be used for matching or other applications.

**Messy and biased data** Many policymakers voiced a lack of trust in the recorded HMIS data (P0, P1, P2, P6, P8). Explicit descriptions of the data included “horrible” (P0), “messy” (P2), “really crappy” (P6), “not reflective of what we’re seeing” (P7), and “inaccurate a lot of times” (P11). Reasons for these characterizations included human error such as case managers not submitting or updating notes or not updating clients’ information within the HMIS (P3, P4).

*“Do we have total faith that all the data on an individual is accurate? If I were to build an AI model that would say who is the best match for it [PSH], you have to have some level of trust that whatever data gets fed into the AI learning model is good data. And I worry about that a little bit.” (P2)*

Policymakers mentioned a lack of follow-up in client outcomes after allocating them a resource, details vital to building confidence in an AI matching system that utilizes machine learning assisted causal inference techniques (P6, P8):

*“We’ll connect someone somewhere and then we’ll say ‘Check! Housed.’ But there’s not a lot of data that says ‘Was this good?’ Even though we know that there are people who transition from one housing place to another housing place because they need to level up or down. I don’t know that that is tracked at all.” (P6)*

There is worry that any AI system, whether used for housing decisions or other applications, would perpetuate biased data, as many policymakers directly referenced the evidence of racial bias in the VI-SPDAT (P1, P2, P5, P6, P9, P11):

*“AI is only as smart as the data [...] so I think there’s the potential for it [AI] to recreate a lot of challenges, [...] a lot of potentially ingrained racist policies of the past.” (P6)*

Beyond the VI-SPDAT, P0 mentioned existing racial biases in who receives services through street outreach programs, where Hispanic individuals receive services at lower rates than they are experiencing homelessness: *“A lot of times we want to focus on the matching and the housing aspect of it, but that is only informed by who housing navigation is actually delivering services to.”* Thus, there exists a racial bias in who is (or is not) even represented within the data.

Policymakers noted bias due to human elements of housing. For example, newer case managers may lack training in recording data or administering the VI-SPDAT (P1, P3, P8). Or perhaps a client *“gets lucky”* with a better case worker (P2, P11). Various policymakers noted a bias since data (e.g., the VI-SPAT, case notes, etc.) is collected through a *“human-interaction”* process and can be subjective or misinterpreted (P2, P3, P8, P11). There is also doubt that the recorded VI-SPDAT answers are accurate. Clients may know that a higher score leads to housing, and case managers want to get their clients housed, leading to potentially *“fluffed up”* answers (P4) and inflated VI-SPDAT scores (P4, P6, P12).

**Lack of data practices** Policymakers mentioned a culture that does not require or encourage standardized data practices (P1, P4). In particular, P1 noted that this is highly variable from agency to agency. Thus, if an AI matching system were introduced, this could lead to even more bias due to disparities in data quality:

*“Some of the agencies that don’t look at HMIS as a positive thing, their participants would lose out. Because we have small agencies, big agencies, grass-roots [...] and it would have to be this universal agreement that this [data quality] is the priority, because we’re going to be getting these [prioritization] lists [...] And so if you didn’t do it right, your people might miss out.” (P1)*

P1 noted the lack of attention to data may be a result of social work culture, as *“social workers aren’t great at math [...] so when you’re asking ‘people people’ that only want to interface with people, data falls to the bottom.”* However, they noted that their particular agency has had a stronger push toward data practices within the past few years.

**AI Tools to Improve Data Processes** Beyond AI for matching, policymakers envisioned various other AI applications that could assist in the data challenges (P6, P7, P8,

P12). For example, P7 mentioned that their organization is using AI to *“make data more accessible for providers.”* P12, in discussing that answers to the VI-SPDAT may not always be truthful, suggested a way of using AI to populate certain information based on verifiable data.

*“I honestly don’t have full understanding of how LAHSA manages their data right now, but it doesn’t feel nimble [which] makes it difficult to make decisions and to gauge whether equitable outcomes exist [...] So any way that AI could help, even if not with the decision-making part of it, but the way that we are able to work with data [...] I think that that could be really huge.” (P8)*

## **Perceptions of AI: the Good, the Bad, and the Unsure**

In general, with the discussion of the case study of Tang et al. (2023), policymakers were excited about the possibilities of incorporating AI tools in the matching process to improve efficiency and remove human bias. They also voiced concern and doubt in how AI could address or be integrated into various aspects of the housing system. As to the specifics of the case study, policymakers valued the proposed equity gains but did not arrive at a clear consensus as to the merits and practicality of the proposed transparency aspects. We first discuss policymakers’ thoughts related to incorporating any AI tool into the matching system and then their thoughts on the specifics of the Tang et al. (2023) case study.

**Benefits of Using AI for Matching** Policymakers viewed AI as a tool that could reduce the amount of tedious work often required of overburdened caseworkers with high turnover rates (P0, P3, P4, P6, P9, P10, P11):

*“I honestly think that [AI] can maybe help improve the epidemic that we’re having right now [...] It can help [matchers] be less stressed out and actually value what they’re doing [...] From what I see is they’re overworked or overtired, and to have a tool that can help them do what they need to do would be very beneficial.” (P9)*

P12 noted the potential benefit of efficiency, where *“[with AI] it takes [fewer] people to figure out things, which is maybe replacing jobs. But maybe [these are] jobs that are hard to fill and people aren’t staying in [...] It could potentially help the system move things faster.”* In general, many policymakers thought AI could help those in homeless services be more “efficient” or to make parts of the housing process “simpler” (P0, P2, P3, P6, P7, P8, P10, P11, P12).

Policymakers voiced that they were hopeful AI could reduce some human bias in the matching process. Though they recognized that humans add bias in various ways, they still want humans involved, likening AI to an extra resource to consider as part of the matching process (P0, P1, P2, P3, P4, P6, P7, P9, P11, P12). As P11 noted, *“because the work we’re doing is human service work [...] I would hate to see it completely gone, the human aspects to the mix.”*

**Risks of using AI for matching** Policymakers had various doubts that AI could replace all parts of the housing

process, even though this is not necessarily the intent. For example, P1 noted that sometimes a particular SPA will not have enough PSH units: “*You have to go and beg other SPAs, ‘Can we please get some units?’ And that again is one of those things that AI can’t [do].*” P5 strongly opposed the use of AI for matching clients, noting “*I just can’t see a computer doing something that is supposed to be a human connection. And it’s gonna sever that.*”

There was a worry of becoming reliant on any AI system, noting that a tool is only effective if it is actually embraced by the end users. Thus, it is important to consider the larger decision-making environment (P6, P9):

*“People tend to follow them [guidelines] because they think they don’t know all the history. And so, you might find that it’s a suggestion or a preference that becomes a rule [...] So I think that’s one of the dangers ... [Case managers are] making these judgments in a very short window of time with a lot of stress. So it’s also the culture of who’s utilizing the output that can also complicate.”* (P6)

**Divergent views on equity** In general, policymakers appreciated the potential equity gains under the system of Tang et al. (2023) as shown in Figure 1 (P2, P3, P4, P10, P12). They did not agree as to whether the AI system should focus on outcome parity or allocation parity as the preferred metric. For example, P0 noted that historical biases in the data could skew the outcome parity model: “*What if historically we’ve just given resources to White and Black folks, and not Latin folks or Asian folks? So then they get the short end of the stick because ... [AI is] building off of whatever happened historically.*” Instead, they “*felt better*” about the allocation parity model (P0). Meanwhile, when discussing the allocation parity model, P7 stated “*not every person needs permanent supportive housing [...] I think that when we create AI, [AI developers] have to be able to understand [that].*”

**Challenges of operationalizing transparency** Policymakers did not agree, even within themselves, as to whether they saw the transparency aspects of the AI matching system of Tang et al. (2023) as an improvement over the current system. Even if they did appreciate such aspects, they grappled with how the proposed system might operationalize in practice (P1, P2, P3, P6, P8, P11, P12). For example, P1 noted “*I would love it if we could say, ‘Look, you’re this place in line, and it’s estimated this time.’ I just don’t understand how we can do that in our system, the way it’s built right now.*” P12 shared similar mixed sentiments, but recognized the needs of clients versus their own sentiments:

*“I’m not certain how I feel about it [clients knowing their place in line]. However, if I were homeless, then I would want to know when might I have a housing resource? I can see why people would want to know that. But I also think it’s somewhat dangerous because there are so many factors [...] The system is very complex.”* (P12).

Stakeholders compared the proposed queue model for housing resources to settings that utilize physical queues such as a “*deli*” (P2) or a “*supermarket*” (P6). They noted

however, these models may not translate well to the housing setting. For example, P2 stated that one can get an estimated wait time for an order at the deli but in housing “*you’re talking about something that’s much higher stakes than a sandwich.*” P6 additionally pointed out that it is easier to implement queues where there is a physical line instead of in the housing system where clients are in line “*on paper.*”

Policymakers noted concerns in providing particular information, such as estimated wait times, to clients directly, as opposed to their case worker. For example, if the housing timeline does not pan out, for whatever reason, this leads to disappointing clients and “*breaking their trust*” (P0). However, we note that sharing this information with clients is contentious within the current housing system (P0, P1, P2, P6, P10). As P6 explained, clients will “*be less likely to want to engage with the government if [they] knew that it would take [them] 10 years to get anything.*” Similarly, P10 voiced that if estimated wait times are shared with clients, they may become less likely to fully engage in their own housing stabilization efforts. For example, they may not feel compelled to continue meeting with their mental health worker since they are “*gonna be next [for housing] really soon*” (P10).

## Discussion

Our study engaged with policymakers in homeless services in LA, discussing the current challenges of the matching process and their openness and hesitations to using AI to assist in this process. Using the AI matching tool of Tang et al. (2023) as a case study, we presented potential efficiency, equity, and transparency features that could be achieved through an AI matching system, gaining feedback on any value policymakers saw in these. We also discussed potential worries about using AI within the matching process. We now discuss the implications of these findings.

### Adopting AI Tools in Homeless Services

Our findings show that even with the awareness of data issues and potential risks to using AI, policymakers are willing to engage with and explore AI tools that are thoughtfully designed and implemented to improve the matching process for both service providers (matchers and case workers) and clients. As seen in Table 2, by the end of the interview, all but one policymaker recognized that AI could provide some positive change in the matching process. This sentiment holds even if all the implementation details are not finalized, as seen within our case study, demonstrating an openness to explore the idea in general.

Though the interview questions were mainly designed to discuss AI within the matching process and the case study of Tang et al. (2023), policymakers offered insights into different applications within or beyond matching where they envisioned AI assisting in homeless services. This provides multiple avenues of exploration for collaboration between AI researchers, developers, and those in homeless services to create such tools and assistants.

However, to build an AI matching tool that is readily adopted by the homeless services community, AI researchers must address multiple issues. We believe these issues are relevant to public services in general.

**Improving data quality and practices** There is currently an effort in LA to both utilize an updated version of the VI-SPDAT that more accurately assesses vulnerability and leverage additional, relevant data sets (i.e., county health data) (Rice et al. 2023). Previous AI research has also shown that messy or missing data can potentially be corrected or accounted for, see e.g., Justin et al. (2023); Bertsimas, Delarue, and Pauphilet (2024); Kern, Kim, and Zhou (2024).

Beyond these potential avenues, there are no standardized data practices across service providers in LA. With many competing day-to-day priorities and a lack of data culture as mentioned above, it may be difficult for those in public services to understand and justify dedicating resources to this effort. By creating proof-of-concept AI tools, AI researchers can show the “value” of prioritizing data in public services to policymakers who have the ability to affect such changes.

To create more faithful data for an AI matching system and to reduce potential biases across providers with different data standards, the adoption of uniform data practices is imperative. AI researchers must collaborate with service providers to define and integrate data quality standards and practices, potentially changing the culture within providers to prioritize these habits, as was previously discussed. For example, case workers can be trained on how to most accurately input client data and “best practices” if doing so is not feasible. In particular, unobserved confounders are of concern for causal inference methods such as Tang et al. (2023), so it is important that all information used in the matching process is recorded. One potential avenue for exploration is working closely with agencies that have been more successful at integrating data practices into their workflow.

**Building trust through evaluation** Even when thoughtfully designed, AI systems can perform differently when deployed due to issues such as shifting distributions of data (Justin et al. 2023; Kern, Kim, and Zhou 2024), unobserved confounders (Kallus and Zhou 2021), and changing system requirements as seen with LAHSA. To build community trust in an AI matching system, AI researchers must periodically audit its performance to measure if it performs as desired by service providers and calibrate its performance as needed. From a data perspective, this requires a commitment to long-term follow up with clients (e.g., more than 2 years after receiving or not receiving a resource) to measure whether clients are staying housed with their given resource.

From an operational perspective, evaluation should also involve follow-up (e.g., interviews) with how matchers engage with such an AI tool. This includes determining how often they follow the suggested matches compared to their on-the-ground expertise. Ideally, when matchers make decisions that differ from the AI system, this information can be used to further refine or train the model, moving towards a more useful tool for matchers.

**Multi-stakeholder engagement and agreement** Policymakers did not agree on what the AI system of Tang et al. (2023) should achieve. For example, there was no consensus in whether the AI matching policy should achieve allocation or outcome parity, or whether there are merits to using a first come, first served queue model. AI researchers must

continue to engage with the community to determine the ultimate, desired properties of an AI matching system, even beyond the properties presented here. Though there is rich research in this area across public services, particular communities may desire different algorithmic design principles, thus, warranting this investigation (Lee et al. 2019; Johnston, Blessenohl, and Vayanos 2020; Stapleton et al. 2022; Johnston et al. 2023; Kuo et al. 2023). This could involve additional input from current or former clients in addition to those who are policymakers and service providers.

## Limitations and Future Work

Policymakers overwhelmingly voiced that if AI is integrated into the matching process, a human decision-maker should still be involved. Thus, a human matcher can use an AI matching tool as a suggestion or can “override” the tool. However, the results of Figure 1 are shown in regards to a decision-maker who follows the AI recommendations faithfully. Thus, it remains to be seen what efficiency, equity, and transparency gains are achieved with the integration of human decision-making and AI and how AI can further learn from the human decision-making process.

Another limitation of this work is that a minority of interviewees have lived homelessness experience (see Table 1). It is important to continue to engage with and educate individuals in this community on the potential applications, merits, and downsides of using AI within the matching process and homeless services in general since this is the population whose livelihoods would be most affected by such a tool. This invites the study and exploration of new methods to engage with the community such as those proposed in Kuo et al. (2023). A related area to explore is that certain sub-populations of clients, such as victims of domestic violence, need additional consideration around issues of privacy and confidentiality if using their personal data in an AI system.

As noted previously, though our study participants are experts in homeless services in LA, issues of data quality, human bias, and high staff turnover are not unique to LA or even homeless services. Though relevant stakeholders’ willingness to accept data-driven tools must be established in other locations and domains, continued research efforts in building data-driven decision assistance tools that address these challenges are an important step in creating tools that may benefit these communities.

## Conclusion

In this work we conducted semi-structured qualitative interviews with 13 policymakers in homeless services in LA on current issues with the housing resource allocation system and their perceptions of utilizing AI tools in this space. Using the AI allocation policy of Tang et al. (2023) as a case study, we presented potential efficiency, fairness (equity), and transparency features of such a system and received feedback on these topics. We found that policymakers are in general open to the use of AI tools that are thoughtfully designed to assist in the matching process if human-decision makers are still involved. We presented various avenues for AI researchers to explore to engage with experts in homeless services in bringing AI tools into this domain.

## Acknowledgments

The authors gratefully acknowledge Hailey Winetrobe Nadel for her assistance and the Social Work Ph.D. students of the USC Center for AI in Society for their feedback on the interview protocol. They would also like to thank Marina Flores, Bryan Paz, and Paola Casillas for their help in subject recruitment and the participants of the study for their time and insight. P. Vayanos is funded in part by the National Science Foundation under CAREER award number 2046230. C.M. Johnston is funded through the Boucher Charitable Gift Fund.

## Positionality

Our research team is an interdisciplinary mix of academics in Social Work, HCI, AI, and Operations Research (OR). Most of the team has previous research experience and community engagement experience with individuals in homeless services or those with lived homelessness experience. The lead researcher is an OR and AI specialist who conducted the semi-structured interviews and participated in the coding process. Additionally, one of the social work researchers with experience with individuals experiencing homelessness participated in the coding process. The principal investigator is an expert in OR and AI with expertise in advancing AI/OR in directions that advance social justice and support the most vulnerable and under-served communities. They have been collaborating with social work researchers with expertise in homelessness and policymakers in homeless services for a decade. Three of the researchers identify as White women, one identifies as a White and Hispanic/Latino woman, one identifies as an Asian woman, and one identifies as a White man.

## References

Ahani, N.; Andersson, T.; Martinello, A.; Teytelboym, A.; and Trapp, A. C. 2021. Placement Optimization in Refugee Resettlement. *Operations Research*, 69(5): 1468–1486.

Akpinar, N.-J.; Lipton, Z. C.; and Chouldechova, A. 2024. The Impact of Differential Feature Under-reporting on Algorithmic Fairness. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1355–1382.

Azizi, M. J.; Vayanos, P.; Eric, R.; and Milind, T. 2018. Designing Fair, Efficient, and Interpretable Policies for Prioritizing Homeless Youth for Housing Resources. In van Hoeve, W.-J., ed., *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 35–51. Springer International Publishing.

Bertsimas, D.; Delarue, A.; and Pauphilet, J. 2024. Simple Imputation Rules for Prediction with Missing Data: Theoretical Guarantees vs. Empirical Performance. *Transactions on Machine Learning Research*.

Brayne, S.; and Christin, A. 2021. Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts. *Social Problems*, 68: 608–624.

Brown, M.; Cummings, C.; Lyons, J.; Carrion, A.; and Watson, D. 2018. Reliability and validity of the Vulnerability

Index-Service Prioritization Decision Assistance Tool (VI-SPDAT) in real-world implementation. *Journal of Social Distress and the Homeless*, 27.

Chouldechova, A.; Benavides-Prado, D.; Fialko, O.; and Vaithianathan, R. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 134–148. PMLR.

City of Hamilton. 2024. Encampment Protocol. [https://www.hamilton.ca/sites/default/files/2024-08/Housing\\_Encampment-Protocol-Revised-June2024.pdf](https://www.hamilton.ca/sites/default/files/2024-08/Housing_Encampment-Protocol-Revised-June2024.pdf).

De Jong, I. 2024. A Message from OrgCode on the VI-SPDAT Moving Forward. <https://www.orgcode.com/blog/a-message-from-orgcode-on-the-vi-spdat-moving-forward>.

Denton, J. 2019. Will Algorithmic Tools Help or Harm the Homeless? *Pacific Standard*.

Eubanks, V. 2019. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. Picador.

Flügge, A. A.; Hildebrandt, T.; and Møller, N. H. 2021. Street-Level Algorithms and AI in Bureaucratic Decision-Making. In *Proceedings of the ACM on Human-Computer Interaction*, volume 5.

Fried, I. 2024. OpenAI says ChatGPT usage has doubled since last year. *Axios*.

Homeless Individuals and Families Information System. 2024. Coordinated Access in HIFIS. <https://homelessnesslearninghub.ca/wp-content/uploads/2024/05/HIFISCAGuide-EN-Final.pdf>.

Jo, N.; Tang, B.; Dullerud, K.; Aghaei, S.; Rice, E.; and Vayanos, P. 2023. Fairness in Contextual Resource Allocation Systems: Metrics and Incompatibility Results. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence*.

Johnston, C. M.; Blessenohl, S.; and Vayanos, P. 2020. Preference Elicitation and Aggregation to Aid with Patient Triage during the COVID-19 Pandemic. *Workshop on Participatory Approaches to Machine Learning*.

Johnston, C. M.; Vossler, P.; Blessenohl, S.; and Vayanos, P. 2023. Deploying a Robust Active Preference Elicitation Algorithm on MTurk: Experiment Design, Interface, and Evaluation for COVID-19 Patient Prioritization. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*.

Justin, N.; Aghaei, S.; Gómez, A.; and Vayanos, P. 2023. Learning Optimal Classification Trees Robust to Distribution Shifts. arXiv:2310.17772.

Kallus, N.; and Zhou, A. 2021. Minimax-Optimal Policy Learning Under Unobserved Confounding. *Management Science*, 67: 2870–2890.

Karusala, N.; Wilson, J.; Vayanos, P.; and Rice, E. 2019. The street-level realities of data practices in homeless services provision. In *Proceedings of the ACM on Human-Computer Interaction*, volume 3, CSCW.

- Kern, C.; Kim, M.; and Zhou, A. 2024. Multi-Accurate CATE is Robust to Unknown Covariate Shifts. *Transactions on Machine Learning Research*.
- KPMG. 2024. Current State Assessment Report: United Way of Greater Los Angeles Homeless Sector Recruitment Analysis. Technical report, Homelessness Research Institute, National Alliance to End Homelessness.
- Kube, A. R.; Das, S.; and Fowler, P. J. 2019. Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 622–629.
- Kube, A. R.; Das, S.; and Fowler, P. J. 2023. Fair and Efficient Allocation of Scarce Resources Based on Predicted Outcomes: Implications for Homeless Service Delivery. *Journal of Artificial Intelligence Research*, 76(73): 1219–1245.
- Kumar, A.; and Yeoh, W. 2023. Fairness in Scarce Societal Resource Allocation: A Case Study in Homelessness Applications. In *Proceedings of the Workshop on Autonomous Agents for Social Good*.
- Kuo, T.-S.; Shen, H.; Geum, J.; Jones, N.; Hong, J. I.; Zhu, H.; and Holstein, K. 2023. Understanding Frontline Workers’ and Unhoused Individuals’ Perspectives on AI Used in Homeless Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.
- Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; and Procaccia, A. D. 2019. WeBuildAI: Participatory framework for algorithmic governance. In *Proceedings of the ACM on Human-Computer Interaction*, volume 3, 1–35.
- Los Angeles Homeless Services Authority. 2017. Report and Recommendations of the Ad Hoc Committee on Women and Homelessness. <https://cao.lacity.gov/Homeless/hsc20170928c.pdf>.
- Los Angeles Homeless Services Authority. 2018. Report and Recommendations of the Ad Hoc Committee on Black People Experiencing Homelessness. <https://www.lahsa.org/documents?id=2823-report-and-recommendations-of-the-ad-hoc-committee-on-black-people-experiencing-homelessness.pdf>.
- Los Angeles Homeless Services Authority. 2024a. 2024 Greater Los Angeles Homeless Count Results (Long Version). <https://www.lahsa.org/documents?id=8164-2024-greater-los-angeles-homeless-count-results-long-version-.pdf>.
- Los Angeles Homeless Services Authority. 2024b. 2024 Housing Inventory Count. <https://www.lahsa.org/documents?id=8162-2024-housing-inventory-count.xlsx>.
- Los Angeles Homeless Services Authority. 2024c. Coordinated Entry System Assessment/Matching Guidance & How to Administer the CES Assessment (VI-SPDAT). <https://www.lahsa.org/documents?id=7969-ces-guidance-and-vi-spdats-training.pdf>.
- Los Angeles Homeless Services Authority. 2024d. Webinar Centralized Matching Presentation. <https://www.lahsa.org/documents?id=8235-webinar-centralized-matching-presentation.pdf>.
- Mashiat, T.; Gitiaux, X.; Rangwala, H.; Fowler, P.; and Das, S. 2022. Trade-offs between Group Fairness Metrics in Societal Resource Allocation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1095–1105.
- Moon, E. S. Y.; and Guha, S. 2024. A Human-Centered Review of Algorithms in Homelessness Research. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
- Moon, E. S.-Y.; Saxena, D.; Das, D.; and Guha, S. 2025. The Datafication of Care in Public Homelessness Services. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Moses, J. 2023. Working in Homeless Services: A Survey of the Field. Technical report, Homelessness Research Institute, National Alliance to End Homelessness.
- Møller, N. H.; Shklovski, I.; and Hildebrandt, T. T. 2020. Shifting Concepts of Value: Designing Algorithmic Decision-Support Systems for Public Services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*.
- Orgcode. 2014. Vulnerability Index (VI) and Service Prioritization Decision Assistance Tool (SPDAT).
- Padgett, D. 2008. *Qualitative methods in social work research*. Sage Publications, 2nd edition.
- Rahmattalabi, A.; Vayanos, P.; Dullerud, K.; and Rice, E. 2022. Learning Resource Allocation Policies from Observational Data with an Application to Homeless Services Delivery. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1240–1256.
- Reamer, F. G. 2023. Artificial Intelligence in Social Work: Emerging Ethical Issues. *International Journal of Social Work Values and Ethics*, 20.
- Rice, E.; Milburn, N.; Vayanos, P.; Rountree, J.; Hill, C.; Petering, R.; Blackwell, B.; Santillano, R.; Onasch-Vera, L.; Nadel, H. W.; Tang, B.; Aghaei, S.; Hsu, H.; and Petry, L. 2023. CESTTRR: Coordinated Entry System Triage Tool Research and Refinement. Technical report.
- Saxena, D.; Badillo-Urquiola, K.; Wisniewski, P. J.; and Guha, S. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. In *Proceedings of the ACM on Human-Computer Interaction*, volume 5.
- Saxena, D.; and Guha, S. 2024. Algorithmic Harms in Child Welfare: Uncertainties in Practice, Organization, and Street-level Decision-making. *ACM Journal on Responsible Computing*, 1: 1–32.
- Showkat, D.; Smith, A. D.; Lingqing, W.; and To, A. 2023. “Who is the right homeless client?”: Values in Algorithmic Homelessness Service Provision and Machine Learning Research. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.
- Stapleton, L.; Lee, M. H.; Qing, D.; Wright, M.; Chouldechova, A.; Holstein, K.; Wu, Z. S.; and Zhu, H. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In

*Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1162–1177.

Tang, B.; Koçyiğit, Ç.; Rice, E.; and Vayanos, P. 2023. Learning Optimal and Fair Policies for Online Allocation of Scarce Societal Resources from Data Collected in Deployment. arXiv:2311.13765.

Tang, N.; Zhi, J.; Kuo, T. S.; Kainaroi, C.; Northup, J. J.; Holstein, K.; Zhu, H.; Heidari, H.; and Shen, H. 2024. AI Failure Cards: Understanding and Supporting Grassroots Efforts to Mitigate AI Failures in Homeless Services. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 713–732.

The County of Santa Clara. 2024. Coordinated Entry. <https://scc.bitfocus.com/coordinated-entry>.

Toros, H.; and Flaming, D. 2018. Prioritizing Homeless Assistance Using Predictive Algorithms: An Evidence-Based Approach. *Cityscape*, 20(1).

Tracey, P.; and Garcia, P. 2024a. After automation: Homelessness prioritization algorithms and the future of care labor. *Big Data and Society*, 11.

Tracey, P.; and Garcia, P. 2024b. Intermediation: Algorithmic Prioritization in Practice in Homeless Services. In *Proceedings of the ACM on Human-Computer Interaction*, volume 8, CSCW2.

U.S. Department of Housing and Urban Development. 2017. CPD-17-01 Notice Establishing Additional Requirements for a Continuum of Care Centralized or Coordinated Assessment System. Technical report.

Vayanos, P.; Ye, Y.; McElfresh, D. C.; Dickerson, J. P.; and Rice, E. 2020. Robust Active Preference Elicitation. arXiv:2003.01899.

Watson-Daniels, J.; Barocas, S.; Hofman, J. M.; and Chouldechova, A. 2023. Multi-Target Multiplicity: Flexibility and Fairness in Target Specification under Resource Constraints. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 297–311.