

Predictive Modeling of Homeless Service Assignment: A Representation Learning Approach

Khandker Sadia Rahman, Charalampos Chelmis

Department of Computer Science, University at Albany, SUNY
Albany, New York, USA
{srahman2, cchelmis}@albany.edu

Abstract

In recent years, there has been growing interest in leveraging machine learning for homeless service assignment. However, the categorical nature of administrative data recorded for homeless individuals hinders the development of accurate machine learning methods for this task. This work asserts that deriving latent representations of such features, while at the same time leveraging underlying relationships between instances is crucial in algorithmically enhancing the existing assignment decision-making process. Our proposed approach learns temporal and functional relationships between services from historical data, as well as unobserved but relevant relationships between individuals to generate features that significantly improve the prediction of the next service assignment compared to the state-of-the-art.

Code — <https://github.com/IDIASLab/REPLETE>

Extended version — <https://arxiv.org/abs/2412.07747>

1 Introduction

Machine learning has gained significant attention for its ability to solve complex real-world problems across various domains including criminal justice, e-commerce, healthcare, banking, finance, and social service (Sarker 2021). At the same time, a constant rise in the rate of homelessness has been observed (Dej, Gaetz, and Schwan 2020). In the United States alone, homelessness has risen 12.5% between 2022 and 2023, with a particularly sharp increase of 29.5% in chronic patterns of homelessness (Henry et al. 2023). Chronic pattern of homelessness occurs when individuals experience repeated homelessness over an extended period (e.g., at least two years), or endure continuous homelessness for at least 12 months (Fleury et al. 2021).

To date, various machine learning approaches have been proposed for predicting whether individuals will reenter the homeless system (Gao, Das, and Fowler 2017; Hong et al. 2018; Kube, Das, and Fowler 2019) as well as assessing the risk of chronic homelessness of individuals (Van-Berlo et al. 2021; Messier, John, and Malik 2021). However, only limited work has addressed the more challenging problem of homeless service assignment (Rahman and Chelmis

2022; Rahman, Zois, and Chelmis 2023; Pokharel, Das, and Fowler 2024). Specifically, (Pokharel, Das, and Fowler 2024) investigate simple models (e.g., Decision Trees with Short Explainable Rules) to recommend the exact service individuals can benefit from, while (Rahman and Chelmis 2022; Rahman, Zois, and Chelmis 2023) explicitly model the trajectories of individuals within the homeless system. The main limitation of these methods lies in their inability to capture the informative relationships between different feature values (i.e., within-feature interactions) and the complex between-feature interactions.

Administrative data collected by the homeless service providers comprises of (i) service assignments and duration of stay, (ii) housing outcome after service assignments, (iii) demographics (e.g., race, gender, ethnicity), educational history, disabling condition, and (iv) time-variant information (e.g., monthly income, health). Most of these features are categorical, and existing methods use one-hot encoding to handle them, mainly for simplicity. Unfortunately, such treatment misses the rich relationships within the features, while exponentially increasing the dimensionality of the data with growing number of categories (Rodríguez et al. 2018). Instead, we propose a predictive model that learns latent representations for features and services, which are then utilized to capture feature interactions and relationships between individuals. We show that the learned representations significantly improve the accuracy of homeless service assignment compared to the state-of-the-art.

This brings us to the core contribution of our work. Similar to (Rahman and Chelmis 2022; Rahman, Zois, and Chelmis 2023) we study the interactions between services found in individuals’ history. However, those studies oversimplify the task of predicting the likelihood of reaching the next service based solely on transitional probabilities between past services, while we assert that many other interactions exist that provide valuable context which can in turn enhance predictive performance. Specifically, subsequent services are often assigned to individuals to serve a particular purpose at a specific time. Understanding such functional and temporal relationship between services can provide insights into why a particular service is assigned at a specific time. Our work extends beyond previous studies by incorporating temporal and functional interaction between services. Furthermore, our model explores the relationships

between individuals. Generally, individuals with similar features tend to receive similar services, and leveraging such information can significantly improve the robustness of predictive models. While (Rahman and Chelms 2022) incorporates finding the “most similar individual” in their prediction problem, it does so as an information retrieval task; instead our approach clusters individuals and utilizes the collective patterns within clusters to provide better insights and more accurate predictions.

In summary, given the history of services assigned to individuals and their current socio-economic features, we propose an objective function with three distinct components that captures (i) temporal and (ii) functional interactions between services, and (iii) interactions between individuals (instances). Through optimization, we obtain latent representations for services and features, and identify the clusters to which each individual belongs. Next, features derived from these representations are input into a feed-forward neural network for predicting the next service assignment. We state our main contributions as follows:

1. We propose a representation-based predictive model for homeless service assignment that effectively learns the latent representation of services and features.
2. We design a novel optimization function that incorporates temporal, functional, and instance interactions to enhance the learning of representations.
3. We utilize the learnt representations to derive features that significantly improve the performance of service assignment predictions.

The rest of the paper is organized as follows. Section 2 summarizes the related work. Section 3 delineates the problem setting. Section 4 introduces the proposed approach. Section 5 describes the data, metrics, and baselines. Section 6 provides detailed discussion of the experimental results. Section 7 concludes with a discussion of the limitations of our study and potential directions for future work.

2 Related Work

Reentry Prediction and Risk of Chronic Homelessness

A considerable body of prior work has focused on predicting reentry into homelessness and assessing the risk of homelessness or chronic homelessness. (Gao, Das, and Fowler 2017; Hong et al. 2018) investigated the application of machine learning models such as logistic regression and random forest to predict whether individuals will reenter homelessness. On the other hand, (VanBerlo et al. 2021; Messier, John, and Malik 2021) explored neural networks to assess the risk of chronic homelessness. Additionally, (Vajiac et al. 2024) investigated various machine learning models to evaluate the risk of eviction-caused homelessness and accurately identify individuals in need of assistance. All of these works are formulated as binary classification problems. Our work, on the other hand, focuses on developing a predictive model for service assignment task, which is a more challenging multi-class classification problem.

Homeless service assignment

Another body of work focuses on predicting the service as-

signment at various conditions, such as (Toros and Flaming 2018; Shinn et al. 2013; Greer et al. 2016) develop models that prioritize housing for those at high risk of homelessness. In contrast, our work centers on individuals who have already experienced chronic patterns of homelessness. On the other hand, (Chelms et al. 2021) explores machine learning models to predict the exact support corresponding to the assigned service. Similarly, (Pokharel, Das, and Fowler 2024) employs decision trees with simple explainable rules (SER-DT) algorithm, which divides the problem into multiple binary classification tasks using a one-versus-all prediction approach for service assignment upon entry. While these studies use one-hot encoding for categorical features, in contrast, our approach leverages representation learning to uncover the rich relationships within these features, moving beyond one-hot encoding. Our research aligns with (Rahman and Chelms 2022) and (Rahman, Zois, and Chelms 2023) which address the multi-class service assignment task. Specifically, (Rahman and Chelms 2022) infers a homelessness network and (Rahman, Zois, and Chelms 2023) utilizes a Bayesian network to predict the next service assignment. In contrast, our work integrates both history of individuals within the homeless system and socio-economic features to predict the next service assignment. A key distinction is that while previous studies explore simple transitional relationships between services, our approach leverages more complex (temporal and functional) relationships between services, and relevant interactions between individuals to predict the next service assignment.

3 Problem Setting

We denote the set of individuals with chronic patterns of homelessness as $\mathcal{U} = \{u_1, \dots, u_{|\mathcal{U}|}\}$. The set of homeless services, such as permanent housing, rapid rehousing, day shelter, is denoted by $\mathcal{A} = \{a^1, a^2, \dots, a^{|\mathcal{A}|}\}$. For each $u_i \in \mathcal{U}$, let $\mathcal{T}_{u_i} = \{(a_{t_1}, t_1), \dots, (a_{t_N}, t_N)\}$ represent the history of services assigned to individual u_i , where a_{t_i} signifies assignment to a service at time t_i . Next, we denote $\mathbf{X} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{F}|}$ as the feature matrix which comprises the feature set $\mathcal{F} \in \mathbb{R}^d$ of individuals $u_i \in \mathcal{U}$ at time t_{N+1} .

Given matrix \mathbf{X} and the history of prior service assignments $\mathcal{T}_{u_i} = \{(a_{t_1}, t_1), \dots, (a_{t_N}, t_N)\}$ for individuals $u_i \in \mathcal{U}$, the goal of this paper is to accurately predict the next service assignment $a_{t_{N+1}} \in \mathcal{Y}$ for $\forall u_i$.

4 Proposed Approach

In this section, we present a comprehensive overview of REPLETE, a REpresentation Learning-based sERVICE assignment modEL, which consists of two main components: (i) representation learning framework and (ii) prediction model. The representation learning framework is designed to learn the latent representations of services. The prediction model leverages the learned representations to enhance the homeless service assignment decision-making process. Figure 1 provides an overview of REPLETE.

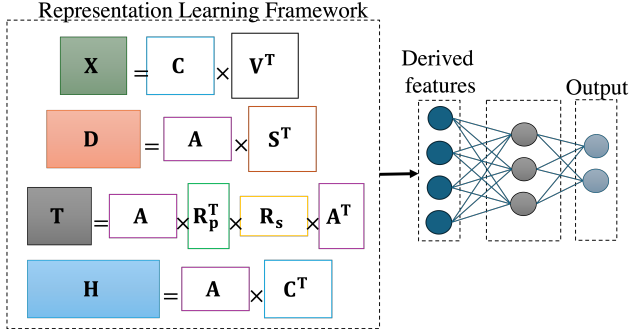


Figure 1: Overview of REPLETE. Representation learning framework first learns representations $\mathbf{A}, \mathbf{C}, \mathbf{V}, \mathbf{S}, \mathbf{R}_p, \mathbf{R}_s$. These are used to derive features, which are subsequently input into FFNN for service assignment prediction.

4.1 Representation Learning Framework

To effectively learn the representations of services, we design an optimization function that captures three different contexts. The first context focuses on the *temporal* interactions between services, where assignments at different stages in the homeless system are likely to have distinct vector representations. The second context captures the *functional* interactions between services, where assignments serving similar needs share similar vector representations. The last context models relationships between *individuals*, where similar individuals are more likely to exhibit similar traits.

Temporal Context Here, our goal is to learn the latent representations of services and time units based on the intuitive observation that services occurring within similar time units should have analogous vector representations. To formalize this, we define $\mathbf{D} \in \mathbb{R}^{|\mathcal{A}| \times |\tau|}$, where \mathbf{D}_{ij} represents the frequency with which service a^i appears in time unit τ_j . Here, τ denotes the set of time units (e.g., weeks, months, or years). To derive standardized representations, we employ Non-negative Matrix Factorization (NMF). Specifically, given \mathbf{D} , we learn the non-negative matrices, $\mathbf{A} \in \mathbb{R}^{|\mathcal{A}| \times k}$ for services and $\mathbf{S} \in \mathbb{R}^{|\tau| \times k}$ for time units by solving

$$\min_{\mathbf{A}, \mathbf{S} \geq 0} \|\mathbf{D} - \mathbf{A}\mathbf{S}^T\|_F^2 + \alpha \sum_{j=1}^{|\tau|-1} \|\mathbf{S}_{j,:} - \mathbf{S}_{j-1,:}\|_F^2 + \lambda(\|\mathbf{A}\|_F^2 + \|\mathbf{S}\|_F^2),$$

where α, λ , and k control the smoothness, sparsity constraints, and dimension of representations, respectively.

Functional Context This context is based on the premise that subsequent services are assigned to individuals because they address needs unmet by the preceding service, meaning, the two services together complement each other in satisfying a need. In this context, we capture the latent relationships between such service pairs. We define a transitional probability matrix $\mathbf{T} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}$, where each element \mathbf{T}_{ij}

denotes the probability of assigning service a^j after service a^i (Rahman and Chelms 2022). Consequently, we learn m latent relations between services represented by matrices, $\mathbf{R}_p \in \mathbb{R}^{m \times k}$ for preceding assignments and $\mathbf{R}_s \in \mathbb{R}^{m \times k}$ for succeeding assignments such that for any pair (a^i, a^j) , $\mathbf{T}_{ij} = \sum_{l=1}^m a^i \mathbf{R}_p[l] + a^j \mathbf{R}_s[l]$. We perform symmetric NMF (Kuang, Ding, and Park 2012) to obtain the solution to

$$\min_{\mathbf{A}, \mathbf{R}_p, \mathbf{R}_s \geq 0} \|\mathbf{T} - \mathbf{A}\mathbf{R}_p^T \mathbf{R}_s \mathbf{A}^T\|_F^2 + \lambda(\|\mathbf{A}\|_F^2 + \|\mathbf{R}_p\|_F^2 + \|\mathbf{R}_s\|_F^2),$$

where \mathbf{A} denotes the latent representations of services. Similar to the temporal context, we incorporate a sparsity constraint with λ controlling the penalty for overfitting, and k denoting the dimension of the representations.

Individual Context In this context, our objective is to learn the latent representations of services denoted by \mathbf{A} and features denoted by $\mathbf{V} \in \mathbb{R}^{|\mathcal{F}| \times k}$, while also identifying clusters to which individuals belong (represented by $\mathbf{C} \in \mathbb{R}^{|\mathcal{U}| \times k}$). By clustering individuals with similar features and past service assignments, our goal is to ensure that services assigned to similar individuals (i.e., in the same cluster) have similar representations. We define $\mathbf{H} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{U}|}$, where each element \mathbf{H}_{ij} denotes the number of times individual u_j is assigned to service a^i . Since \mathbf{X} and \mathbf{H} exhibit higher sparsity, L_{2-1} norm is applied to prevent rows with greater sparsity from dominating the objective function. Given matrices \mathbf{H} and \mathbf{X} , we solve

$$\min_{\mathbf{A}, \mathbf{V} \geq 0, \mathbf{C}} \|\mathbf{H} - \mathbf{A}\mathbf{C}^T\|_{2,1} + \|\mathbf{X} - \mathbf{C}\mathbf{V}^T\|_{2,1} + \lambda(\|\mathbf{A}\|_F^2 + \|\mathbf{V}\|_F^2 + \|\mathbf{C}\|_F^2) + \beta \text{tr}(\mathbf{C}^T \mathbf{\Gamma} \mathbf{C})$$

s.t. $\mathbf{C}\mathbf{1}_k = \mathbf{1}_{|\mathcal{U}|}, \mathbf{C} \in [0, 1]$.

Here, the term $\text{tr}(\mathbf{C}^T \mathbf{\Gamma} \mathbf{C})$ ensures that similar individuals are clustered together and $\mathbf{\Gamma} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{U}|}$ records the cosine similarity between individuals in the set \mathcal{U} . The constraints on \mathbf{C} enforce soft clustering by ensuring that the rules of probability are satisfied. Such soft clustering is appropriate for homelessness, a complex domain where individuals can exhibit similarities with members of different clusters. Finally, β controls the strictness of the clusters by determining how tightly or loosely data instances are grouped together.

Overall Optimization Function With all the previously introduced components, we formulate our joint optimization problem as follows:

$$\min_{\mathbf{A}, \mathbf{S}, \mathbf{V}, \mathbf{R}_p, \mathbf{R}_s \geq 0, \mathbf{C}} \|\mathbf{D} - \mathbf{A}\mathbf{S}^T\|_F^2 + \alpha \sum_{j=1}^{|\tau|-1} \|\mathbf{S}_{j,:} - \mathbf{S}_{j-1,:}\|_F^2 + \|\mathbf{H} - \mathbf{A}\mathbf{C}^T\|_{2,1} + \|\mathbf{X} - \mathbf{C}\mathbf{V}^T\|_{2,1} + \beta \text{tr}(\mathbf{C}^T \mathbf{\Gamma} \mathbf{C}) + \|\mathbf{T} - \mathbf{A}\mathbf{R}_p^T \mathbf{R}_s \mathbf{A}^T\|_F^2 + \lambda(\|\mathbf{A}\|_F^2 + \|\mathbf{S}\|_F^2 + \|\mathbf{V}\|_F^2 + \|\mathbf{C}\|_F^2 + \|\mathbf{R}_p\|_F^2 + \|\mathbf{R}_s\|_F^2)$$

s.t. $\mathbf{C}\mathbf{1}_k = \mathbf{1}_{|\mathcal{U}|}, \mathbf{C} \in [0, 1]$. (1)

4.2 Optimization Algorithm

Jointly updating the variables in Eq. (1) causes the objective function to become non-convex. We therefore optimize the objective function using Alternating Direction Method of Multiplier (ADMM) (Boyd et al. 2011), where variables are updated separately. All proofs are provided in the Appendix within the extended version in arXiv.

We begin by relaxing the constraints on \mathbf{C} to orthogonality, specifically $\mathbf{C}^\top \mathbf{C} = \mathbf{I}$, $\mathbf{C} \geq 0$ (Tang and Liu 2012). We then introduce two auxiliary variables: $\mathbf{P} = \mathbf{H} - \mathbf{A}\mathbf{C}^\top$ and $\mathbf{Q} = \mathbf{X} - \mathbf{C}\mathbf{V}^\top$. Consequently, Eq. (1) is reformulated into the following equivalent problem:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{S}, \mathbf{V}, \mathbf{R}_p, \mathbf{R}_s, \mathbf{C} \geq 0, \mathbf{P}, \mathbf{Q}} & \|\mathbf{D} - \mathbf{A}\mathbf{S}^\top\|_F^2 + \alpha \|\mathbf{S}\mathbf{B}'\|_F^2 + \|\mathbf{P}\|_{2,1} \\ & + \|\mathbf{Q}\|_{2,1} + \beta \text{tr}(\mathbf{C}^\top \mathbf{\Gamma} \mathbf{C}) + \|\mathbf{T} - \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top\|_F^2 \\ & + \lambda (\|\mathbf{A}\|_F^2 + \|\mathbf{S}\|_F^2 + \|\mathbf{V}\|_F^2 + \|\mathbf{C}\|_F^2 + \|\mathbf{R}_p\|_F^2 + \|\mathbf{R}_s\|_F^2) \\ & + \langle \mathcal{L}, \mathbf{X} - \mathbf{C}\mathbf{V}^\top - \mathbf{Q} \rangle + \langle \mathcal{K}, \mathbf{H} - \mathbf{A}\mathbf{C}^\top - \mathbf{P} \rangle \\ & + \langle \mathcal{N}, \mathbf{C}^\top \mathbf{C} - \mathbf{I} \rangle + \frac{\mu}{2} \|\mathbf{H} - \mathbf{A}\mathbf{C}^\top - \mathbf{P}\|_F^2 \\ & + \frac{\mu}{2} \|\mathbf{X} - \mathbf{C}\mathbf{V}^\top - \mathbf{Q}\|_F^2, \end{aligned} \quad (2)$$

where \mathcal{K} , \mathcal{L} , and \mathcal{N} are Lagrangian multipliers, μ controls the penalty for violating the equality constraints, $\langle \cdot, \cdot \rangle$ denotes the dot product, and \mathbf{B}' represents the matrix equivalent to the smoothing constraint. Next, we summarize the update rules for each variable.

Update P For updating \mathbf{P} , we hold the other variables constant and remove the terms that are irrelevant to \mathbf{P} . Eq. (2) can be written as:

$$\min_{\mathbf{P}} \frac{1}{2} \|\mathbf{P} - (\mathbf{H} - \mathbf{A}\mathbf{C}^\top + \frac{1}{\mu} \mathcal{K})\|_F^2 + \frac{1}{\mu} \|\mathbf{P}\|_{2,1} \quad (3)$$

Lemma 1 provides a closed form solution for Eq. (3).

Lemma 1 (Wang, Tang, and Liu 2015) *Given matrix \mathbf{E} and a positive scale α , the i^{th} row of the optimal solution \mathbf{W}^* of $\min_{\mathbf{W}} \frac{1}{2} \|\mathbf{W} - \mathbf{E}\|_F^2 + \alpha \|\mathbf{W}\|_{2,1}$ is given by:*

$$w_i^* = \begin{cases} (1 - \frac{\alpha}{\|e_i\|})e_i, & \|e_i\| > \alpha \\ 0, & \text{otherwise} \end{cases}$$

Using Lemma 1, the optimal solution \mathbf{P}^* for the above equation is as follows, where $\mathbf{E}^P = \mathbf{H} - \mathbf{A}\mathbf{C}^\top + \frac{1}{\mu} \mathcal{K}$:

$$\mathbf{P}_{i,:}^* = \begin{cases} (1 - \frac{1}{\mu \|\mathbf{E}_{i,:}^P\|}) \mathbf{E}_{i,:}^P, & \|\mathbf{E}_{i,:}^P\| > \frac{1}{\mu} \\ 0, & \text{otherwise} \end{cases}$$

Update Q The optimal solution \mathbf{Q}^* is obtained similar to \mathbf{P} , using Lemma 1, where $\mathbf{E}^Q = \mathbf{X} - \mathbf{C}\mathbf{V}^\top + \frac{1}{\mu} \mathcal{L}$:

$$\mathbf{Q}_{i,:}^* = \begin{cases} (1 - \frac{1}{\mu \|\mathbf{E}_{i,:}^Q\|}) \mathbf{E}_{i,:}^Q, & \|\mathbf{E}_{i,:}^Q\| > \frac{1}{\mu} \\ 0, & \text{otherwise} \end{cases}$$

Update S Let ψ_S be the Lagrangian multiplier for $\mathbf{S} \geq 0$, the Lagrangian function related to \mathbf{S} is, $\mathcal{O}_S = \min \|\mathbf{D} - \mathbf{A}\mathbf{S}^\top\|_F^2 + \lambda \|\mathbf{S}\|_F^2 + \alpha \|\mathbf{S}\mathbf{B}'\|_F^2 - \text{tr}(\psi_S \mathbf{S}^\top)$. The partial derivative of \mathcal{O}_S is $\frac{1}{2} \frac{d\mathcal{O}_S}{d\mathbf{S}} = -(\mathbf{D} - \mathbf{A}\mathbf{S}^\top)^\top \mathbf{A} + \alpha (\mathbf{S}\mathbf{B}')\mathbf{B}' + \lambda \mathbf{S} - \psi_S$. Using the Karush–KuhnTucker (KKT) complementary condition (Boyd and Vandenberghe 2004), i.e., $\psi_S(i, j) \mathbf{S}_{ij} = 0$, we get $\mathbf{S}_{ij} \leftarrow \mathbf{S}_{ij} \frac{\hat{\mathbf{S}}_{ij}}{\tilde{\mathbf{S}}_{ij}}$ (Zhang and Zhang 2017), where

$$\begin{aligned} \hat{\mathbf{S}}_{ij} &= \mathbf{D}^\top \mathbf{A} + [\alpha (\mathbf{S}(\mathbf{B}')^2)]^- \\ \tilde{\mathbf{S}}_{ij} &= \mathbf{S}\mathbf{A}^\top \mathbf{A} + \lambda \mathbf{S} + [\alpha (\mathbf{S}(\mathbf{B}')^2)]^+. \end{aligned}$$

Here, for any matrix \mathbf{M} , $(\mathbf{M})^+ = \frac{ABS(\mathbf{M}) + \mathbf{M}}{2}$ and $(\mathbf{M})^- = \frac{ABS(\mathbf{M}) - \mathbf{M}}{2}$ are the positive and negative part of \mathbf{M} , respectively, and $ABS(\mathbf{M})$ consists of the absolute value of elements in \mathbf{M} (Shu, Wang, and Liu 2019).

Update \mathbf{R}_p and \mathbf{R}_s The partial derivative of the Lagrangian objective function w.r.t. \mathbf{R}_p is $\frac{1}{2} \frac{d\mathcal{O}_{\mathbf{R}_p}}{d\mathbf{R}_p} = -\mathbf{T}\mathbf{A}\mathbf{A}^\top \mathbf{R}_s + \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{A}^\top \mathbf{R}_s + \lambda \mathbf{R}_p - \psi_{\mathbf{R}_p}$. Using the KKT complementary condition, we get $\mathbf{R}_{p_{ij}} \leftarrow \mathbf{R}_{p_{ij}} \frac{\hat{\mathbf{R}}_{p_{ij}}}{\tilde{\mathbf{R}}_{p_{ij}}}$, where

$$\begin{aligned} \hat{\mathbf{R}}_{p_{ij}} &= \mathbf{T}\mathbf{A}\mathbf{A}^\top \mathbf{R}_s + (\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{A}^\top \mathbf{R}_s)^- \\ \tilde{\mathbf{R}}_{p_{ij}} &= \lambda \mathbf{R}_p + (\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{A}^\top \mathbf{R}_s)^+. \end{aligned}$$

Similar to \mathbf{R}_p , we get $\mathbf{R}_{s_{ij}} \leftarrow \mathbf{R}_{s_{ij}} \frac{\hat{\mathbf{R}}_{s_{ij}}}{\tilde{\mathbf{R}}_{s_{ij}}}$ for \mathbf{R}_s , such that:

$$\begin{aligned} \hat{\mathbf{R}}_{s_{ij}} &= \mathbf{T}\mathbf{A}\mathbf{A}^\top \mathbf{R}_p + (\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{A}^\top \mathbf{R}_p)^- \\ \tilde{\mathbf{R}}_{s_{ij}} &= \lambda \mathbf{R}_s + (\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{A}^\top \mathbf{R}_p)^+. \end{aligned}$$

Update V The partial derivative of the Lagrangian objective function w.r.t. \mathbf{V} is $\frac{d\mathcal{O}_V}{d\mathbf{V}} = -\mu (\mathbf{X} - \mathbf{C}\mathbf{V}^\top - \mathbf{Q})^\top \mathbf{C} - \mathcal{L}^\top \mathbf{C} + 2\lambda \mathbf{V} - \psi_V$. Using the KKT complementary condition, we get $\mathbf{V}_{ij} \leftarrow \mathbf{V}_{ij} \frac{\hat{\mathbf{V}}_{ij}}{\tilde{\mathbf{V}}_{ij}}$, where

$$\begin{aligned} \hat{\mathbf{V}}_{ij} &= \mu \mathbf{X}^\top \mathbf{C} + (\mathcal{L}^\top \mathbf{C})^+ \\ \tilde{\mathbf{V}}_{ij} &= \mu \mathbf{V}\mathbf{C}^\top \mathbf{C} + \mu \mathbf{Q}^\top \mathbf{C} + (\mathcal{L}^\top \mathbf{C})^- + 2\lambda \mathbf{V}. \end{aligned}$$

Update C The partial derivative of the Lagrangian objective function w.r.t. \mathbf{C} is $\frac{1}{2} \frac{d\mathcal{O}_C}{d\mathbf{C}} = -\mu (\mathbf{H} - \mathbf{A}\mathbf{C}^\top - \mathbf{P})^\top \mathbf{A} - \mu (\mathbf{X} - \mathbf{C}\mathbf{V}^\top - \mathbf{Q})\mathbf{V} - \mathcal{L}\mathbf{V} - \mathcal{K}^\top \mathbf{A} + 2\mathcal{C}\mathcal{N} + 2\beta \mathbf{\Gamma} \mathbf{C} + 2\lambda \mathbf{C} - \psi_C$. Using the KKT complementary condition, we get $\mathbf{C}_{ij} \leftarrow \mathbf{C}_{ij} \frac{\hat{\mathbf{C}}_{ij}}{\tilde{\mathbf{C}}_{ij}}$, where

$$\begin{aligned} \hat{\mathbf{C}}_{ij} &= \mu \mathbf{H}^\top \mathbf{A} + \mu \mathbf{X}\mathbf{V} + (\mathcal{L}\mathbf{V})^+ + (\mathcal{K}^\top \mathbf{A})^+ + 2(\mathcal{C}\mathcal{N})^- \\ &+ 2(\beta \mathbf{\Gamma} \mathbf{C})^- \\ \tilde{\mathbf{C}}_{ij} &= \mu \mathbf{C}\mathbf{A}^\top \mathbf{A} + \mu \mathbf{P}^\top \mathbf{A} + \mu \mathbf{C}\mathbf{V}^\top \mathbf{V} + \mu \mathbf{Q}\mathbf{V} + (\mathcal{L}\mathbf{V})^- \\ &+ (\mathcal{K}^\top \mathbf{A})^- + 2(\mathcal{C}\mathcal{N})^+ + 2\beta \mathbf{\Gamma} \mathbf{C} + 2\lambda \mathbf{C}. \end{aligned}$$

Update A The partial derivative of the Lagrangian objective function w.r.t. \mathbf{A} is $\frac{1}{2} \frac{d\mathcal{O}_A}{d\mathbf{A}} = -2(\mathbf{D} - \mathbf{A}\mathbf{S}^\top)\mathbf{S} - 4(\mathbf{T} - \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top) \times \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s + 2\lambda\mathbf{A} - \mu(\mathbf{H} - \mathbf{A}\mathbf{C}^\top - \mathbf{P})\mathbf{C} - \mathbf{K}\mathbf{C} - \psi_A$. Using the KKT complementary conditions, we get $\mathbf{A}_{ij} \leftarrow \mathbf{A}_{ij} \frac{\hat{\mathbf{A}}_{ij}}{\mathbf{A}_{ij}}$ where

$$\begin{aligned} \hat{\mathbf{A}}_{ij} &= 2\mathbf{D}\mathbf{S} + 4\mathbf{T}\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s + 4(\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s)^- \\ &\quad + \mu\mathbf{H}\mathbf{C} + (\mathbf{K}\mathbf{C})^+ \\ \tilde{\mathbf{A}}_{ij} &= 2\mathbf{A}\mathbf{S}^\top \mathbf{S} + 4(\mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}^\top \mathbf{A}\mathbf{R}_p^\top \mathbf{R}_s)^+ + 2\lambda\mathbf{A} \\ &\quad + \mu\mathbf{A}\mathbf{C}^\top \mathbf{C} + \mu\mathbf{P}\mathbf{C} + (\mathbf{K}\mathbf{C})^-. \end{aligned}$$

Update \mathcal{L} , \mathcal{M} , and \mathcal{N} After updating the variables, we update the Lagrangian parameters (Wang, Tang, and Liu 2015) as follows:

$$\begin{aligned} \mathcal{L} &= \mathcal{L} + \mu(\mathbf{X} - \mathbf{C}\mathbf{V}^\top - \mathbf{Q}) \\ \mathcal{M} &= \mathcal{M} + \mu(\mathbf{H} - \mathbf{A}\mathbf{C}^\top - \mathbf{P}) \\ \mathcal{N} &= \mathcal{N} + \mu(\mathbf{C}^\top \mathbf{C} - \mathbf{I}). \end{aligned}$$

4.3 Service Assignment Prediction

In this section, we delve into the specifics of the service assignment module of REPLETE. Once the representations are obtained by solving Eq. (1), we derive three sets of features: (i) service and feature representations, (ii) feature interactions, and (iii) instance interactions. The overall process is summarized in Algorithm 1.

Service and Feature Representations We utilize the representations of services and features in our prediction model instead of one-hot encoding them. The rationale is that by treating services and features as categorical our model would miss the inherent relationships within them. Therefore, we replace the services $\{a_{t_1}, \dots, a_{t_N}\}$ within the history of services received by an individual, with their respective representations $\{\mathbf{A}_{a_{t_1}}, \dots, \mathbf{A}_{a_{t_N}}\}$, where $\mathbf{A}_{a_{t_i}}$ denotes the representation of service a_{t_i} . We additionally obtain the representations of the socio-economic features for each individual by multiplying \mathbf{X} and \mathbf{V} , where \mathbf{X} is the feature matrix and \mathbf{V} is the matrix of learned feature representations.

Feature Interactions Beyond the representations of services and features, we further incorporate two categories of service interactions: temporal and functional (Section 4.1). Given the sequence of services $\{a_{t_1}, \dots, a_{t_N}\}$ received by individuals, the temporal and functional interactions within each service pair (a_{t_i}, a_{t_j}) are defined as $\mathbf{A}_{a_{t_i}} \mathbf{S}^\top \mathbf{S} \mathbf{A}_{a_{t_j}}^\top$ and $\mathbf{A}_{a_{t_i}} \mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}_{a_{t_j}}^\top$, respectively. Matrix $\mathbf{A}_{a_{t_i}}$ denotes the representation of service a_{t_i} , whereas matrices \mathbf{S} , \mathbf{R}_p , and \mathbf{R}_s are the learned representations for time units, preceding assignments, and succeeding assignments, respectively.

Instance Interactions We define instance interaction as the similarity between individuals, which we use to identify clusters within the individuals. Since the relaxed optimization on \mathbf{C} (i.e., the matrix denoting the clusters to which individuals belong) does not directly provide probabilities for soft clustering, we normalize \mathbf{C} such that each row sums to 1, thereby obtaining these probabilities.

Algorithm 1: REPLETE

Require: $\mathbf{X}, \mathbf{D}, \mathbf{T}, \mathbf{H}, \mathbf{B}', \alpha, \beta, \mu, \lambda, \tau, k$

Ensure: $\hat{a}_{t_{N+1}}$

- 1: Randomly initialize $\mathbf{A}, \mathbf{S}, \mathbf{C}, \mathbf{V}, \mathbf{R}_p, \mathbf{R}_s, \mathbf{P}, \mathbf{Q}$
- 2: Pre-compute similarity matrix Γ
- 3: **repeat**
- 4: Update $\mathbf{A}, \mathbf{S}, \mathbf{C}, \mathbf{V}, \mathbf{R}_p, \mathbf{R}_s, \mathbf{P}, \mathbf{Q}$ according to Section 4.2
- 5: **until** convergence
- 6: Normalize \mathbf{C} s.t. $\sum_j C_{i,j} = 1$
- 7: **for each** $u_i \in \mathcal{U}$ **do**
- 8: $\mathbf{Z}_{u_i} = \mathbf{Z}_{u_i} \cup (\mathbf{X}\mathbf{V})_{u_i} \cup \mathbf{C}_{u_i}$
- 9: **for each** $a_{t_k} \in \mathcal{T}_{u_i}$ **do**
- 10: $\mathbf{Z}_{u_i} = \mathbf{Z}_{u_i} \cup \mathbf{A}_{a_{t_k}}$
- 11: **end for**
- 12: **for each** $(a_{t_m}, a_{t_n}) \in \mathcal{T}_{u_i}$ **do**
- 13: $\mathbf{Z}_{u_i} = \mathbf{Z}_{u_i} \cup (\mathbf{A}_{a_{t_m}} \mathbf{S}^\top \mathbf{S} \mathbf{A}_{a_{t_n}}^\top)$
- 14: $\mathbf{Z}_{u_i} = \mathbf{Z}_{u_i} \cup (\mathbf{A}_{a_{t_m}} \mathbf{R}_p^\top \mathbf{R}_s \mathbf{A}_{a_{t_n}}^\top)$
- 15: **end for**
- 16: **end for**
- 17: Split \mathbf{Z} into training \mathbf{Z}_{train} and testing set \mathbf{Z}_{test}
- 18: Train FFNN with \mathbf{Z}_{train}
- 19: Input \mathbf{Z}_{test} into FFNN
- 20: $\hat{\mathbf{Y}}_{test} \leftarrow$ Output of FFNN
- 21: **return** $\hat{\mathbf{Y}}_{test}$

Prediction Model We use $\mathbf{Z} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{F}'|}$, to denote the concatenated feature vector, i.e., service and feature representations, feature interactions, and instance interactions, where \mathcal{U} is the set of individuals, and \mathcal{F}' is the set of derived features. For the service assignment task, we utilize a single hidden layer feed-forward neural network (FFNN). The rationale behind a single hidden layer architecture is that it balances complexity and computational efficiency.

Predicted service assignments, denoted by $\hat{\mathbf{Y}}$, are calculated as $\hat{\mathbf{Y}} = \sigma(\mathbf{Z}\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2$ where, \mathbf{W}_1 and \mathbf{W}_2 are the weight matrices for the hidden and output layers, respectively, \mathbf{b}_1 and \mathbf{b}_2 are the bias vectors for the hidden and output layers, respectively, and $\sigma(\cdot)$ is the ReLU activation function. The model is trained using cross-entropy loss.

5 Experimental Setup

Data Description Our analysis utilizes an anonymized dataset from CARES of NY comprising 18,817 records from 6,011 chronically homeless individuals in the Capital Region of New York State, covering the period from 2012 to 2018, including a total of 9 different homeless services. A complete description of the services is available at (United States Department of Housing and Urban Development 2020). Sequence of services longer than N are sampled using a sliding window with a length of N . For sequences shorter than N , missing values are encoded with 0, resulting in a zero vector representation with a length $|\mathcal{A}|$. After sampling, we denote our dataset as \mathcal{D} and perform a random 70 – 30 split to obtain the training set \mathcal{D}_{train} and testing

Method	Accuracy	Recall	Precision	F_1 score
TRACE	0.530	0.188	0.205	0.184
Transformer	0.632	0.368	0.406	0.376
LR	0.691	0.406	0.446	0.420
FFNN ₁	0.714	0.385	0.458	0.416
RF	0.738	0.391	0.504	0.434
PREVISE	0.745	0.400	0.575	0.435
RNN	0.800	0.494	0.583	0.511
LSTM	0.802	0.496	0.590	0.513
REPLETE	0.832	0.539	0.615	0.567

Table 1: Performance comparison between REPLETE and the baselines. Parameter N is set to be 3.

set \mathcal{D}_{test} , respectively. Finally, we train and evaluate our approach using \mathcal{D}_{train} and \mathcal{D}_{test} respectively.

Baselines We compare REPLETE with the state-of-the-art: TRACE (Rahman and Chelms 2022) and PREVISE (Rahman, Zois, and Chelms 2023). TRACE and PREVISE uses service sequences to predict the next assignment. Additionally, following (Rahman, Zois, and Chelms 2023), we compare our approach with random forest (RF), logistic regression (LR), transformer, feed forward neural network (FFNN₁), recurrent neural network (RNN), and long short-term memory (LSTM), which incorporate one-hot encoded features for prediction.

Evaluation Metrics We evaluate our approach using accuracy, recall, precision, and F_1 score.

Implementation Details All analyses were conducted in Jupyter Notebook using Python 3 on a MacBook Air (M1 chip, 8 cores, 8 GB memory) running macOS Sonoma. The representation learning framework ran for 500 iterations, and the prediction model averaged 20 iterations per fold in 5-fold cross-validation.

6 Results and Analysis

Quantitative Analysis Table 1 shows that REPLETE significantly outperforms the baselines in predicting the next service assignment, both in terms of accuracy and F_1 score. Table 2 confirms that REPLETE significantly outperforms PREVISE (state-of-the-art) across all metrics and LSTM (the best-performing baseline) in accuracy and F-score. This indicates that (i) the representation learning framework effectively captures the rich relationships within the features and instances, and (ii) derived features based on temporal, functional, and individual contexts, effectively carry these relationships to the FFNN. Together, these factors lead to a significant improvement in the performance of our approach.

To analyze the reliability of the representation learning framework, we first examine whether derived features can effectively distinguish between labels (i.e., services) compared to the one-hot encoded features. We utilize the t-distributed Stochastic Neighbor Embedding (t-SNE) method to embed both high-dimensional features into 2-dimensional space. Figure 2 illustrates that distinct clusters for each label are formed using the derived features. This

t-score (p-value)	Accuracy	Recall	Precision	F_1 score
w.r.t. PREVISE	-15.3 (2.26×10^{-11})	-3.52 (0.0052)	-5.37 (8.8×10^{-5})	-4.07 (0.0018)
w.r.t. LSTM	-5.65 (2.3×10^{-5})	-1.99 (0.064)	-1.86 (0.082)	-2.56 (0.020)

Table 2: Statistical significance comparison between REPLETE and the baselines. Parameter N is set to be 3.




Figure 2: Scatterplot of 2-dimensional t-SNE embedding for (a) one-hot encoded features and (b) derived features from learned representations.

indicates that the derived features offer better separability compared to the one-hot encoded features.

Next, we analyze the performance of REPLETE across different history length N . Table 3 shows that our model’s performance remain consistent despite variation in the length of the individual’s history considered for prediction. This indicates the robustness of our approach. For the rest of our experiments, we set $N = 3$, as the model achieves the highest accuracy and F_1 score at this value.

Ablation Study We conducted two ablation studies to evaluate the importance of different model components, focusing on the impact of feature and instance interactions. First, we analyze the impact of feature interactions and instance interactions on model performance. Table 4 shows that service and feature representations have the greatest impact, with feature and instance interactions enhancing performance when coupled with these representations.

Next, we analyse the impact of each component of the optimization function on model performance. As shown in Table 5, each component enhances performance, with the best results achieved by combining all three components.

Hyperparameter Selection The joint optimization function has six hyperparameters, with optimal values being $\lambda = 0.01$, $\alpha = 0.3$, $\beta = 0.7$, $\tau = 52$ weeks, $\mu = 1$, $k = 10$.

Bias Evaluation Although our proposed approach does not explicitly model against bias and unfairness (both critical characteristics of a technological solution to a complex so-

N	Accuracy	Recall	Precision	F_1 score
2	0.828	0.517	0.644	0.547
3	0.832	0.539	0.615	0.567
4	0.809	0.499	0.571	0.517
5	0.805	0.499	0.566	0.516
6	0.811	0.548	0.607	0.561

Table 3: Performance of our approach across different history length N .

inst	rep	feat	Accuracy	F_1 score
✓			0.625	0.208
	✓		0.825	0.526
		✓	0.614	0.144
	✓	✓	0.829	0.547
✓		✓	0.717	0.279
✓	✓		0.829	0.543
✓	✓	✓	0.832	0.567

Table 4: Ablation study for key components of the derived features, namely service and feature representations (rep), instance interactions (inst), and feature interactions (feat).

ind	temp	func	Accuracy	F_1 score
✓			0.798	0.450
	✓		0.789	0.419
		✓	0.714	0.398
	✓	✓	0.801	0.476
✓		✓	0.821	0.534
✓	✓		0.827	0.549
✓	✓	✓	0.832	0.567

Table 5: Ablation study for key components of the optimization function, namely temporal (temp), functional (func), and individual (ind).

cial challenge, such as homelessness service provision), we evaluate it for potential bias across three sensitive attributes, namely race, gender, and ethnicity. We measure bias using (i) demographic parity ($\frac{P(\hat{Y}=a^i|attr=1)}{P(\hat{Y}=a^i|attr=0)}$) and (ii) equal opportunity ($\frac{P(\hat{Y}=a^i|attr=1, Y=a^i)}{P(\hat{Y}=a^i|attr=0, Y=a^i)}$) (Mehrabi et al. 2021), aiming for values within 80% of the group with the highest rate (Pessach and Shmueli 2022). Figure 3 shows that our approach mitigates bias for high-frequency services, such as emergency shelter (denoted by 1) and day shelter (denoted by 11) for majority of the sensitive attributes, but not as much for lower-frequency services, such as transitional housing (denoted by 2) and homelessness prevention (denoted by 12). Moreover, the state-of-the-art method PREVISE exhibits extreme bias for certain services and attributes, where it completely fails to predict some services for specific attributes. This underscores that our approach, REPLETE, demonstrates a significantly lower level of bias compared to PREVISE, ensuring a more equitable service assignment across various attributes and services. In light of this result, we conclude that our approach replicates the ex-

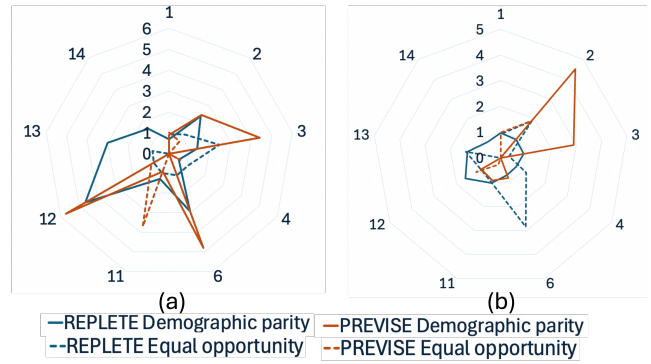


Figure 3: Demographic parity (solid line) and equal opportunity (dotted line) of sensitive attributes (a) gender and (b) ethnicity for each service using REPLETE (blue) and PREVISE (orange). Plots for the remaining attributes are included in the Appendix of the extended version in arXiv. Numerals are used in lieu of actual service names (United States Department of Housing and Urban Development 2020).

isting assignment decision-making process to a significant extent, providing the basis for developing systems that make “unbiased” and “fair” predictions, as well as to understand and evaluate them ethically (i.e., in experimental settings). We therefore plan to incorporate fairness constraints (e.g., (Zafar et al. 2019)) directly into our objective function, as part of our future work.

7 Conclusion

This paper introduced a novel predictive model for homeless service assignment based on representation learning. The proposed approach modeled explicitly both the temporal and functional relationships between services, and the similarity between individuals based on their features and their prior service assignments, to learn latent representations. Utilizing these representations, the proposed approach was shown to outperform the state-of-the-art in the task of next service assignment prediction, a key task in the service assignment decision making process.

Limitations Our analysis is based on a geographically bounded dataset, specifically limited to the Capital Region of New York state. Additionally, the dataset does not record the availability or capacity of services.

Future Directions Our approach serves as a foundational building block for developing more complex service assignment models. While we explored temporal and functional relationships, additional dimensions remain worth investigating, potentially offering valuable insights and new research directions. On the other hand, explicitly modeling against bias and unfairness is necessary, as our experiments showed, particularly if the proposed approach were to one day be used in the real-world. Our work takes the first step in this direction by providing a basis for developing “unbiased” and “fair” predictive models and evaluating them ethically in experimental settings.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. ECCS-1737443.

References

- Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J.; et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1): 1–122.
- Boyd, S.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Chelmiss, C.; Qi, W.; Lee, W.; and Duncan, S. 2021. Smart Homelessness Service Provision with Machine Learning. *Procedia Computer Science*, 185: 9–18.
- Dej, E.; Gaetz, S.; and Schwan, K. 2020. Turning off the tap: a typology for homelessness prevention. *The Journal of Primary Prevention*, 41(5): 397–412.
- Fleury, M.-J.; Grenier, G.; Sabetti, J.; Bertrand, K.; Clément, M.; and Brochu, S. 2021. Met and unmet needs of homeless individuals at different stages of housing reintegration: A mixed-method investigation. *PLoS one*, 16(1): e0245088.
- Gao, Y.; Das, S.; and Fowler, P. 2017. Homelessness service provision: a data science perspective. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- Greer, A. L.; Shinn, M.; Kwon, J.; and Zuiderveen, S. 2016. Targeting services to individuals most likely to enter shelter: Evaluating the efficiency of homelessness prevention. *Social Service Review*, 90(1): 130–155.
- Henry, M.; de Sousa, T.; Roddey, C.; Gayen, S.; Joe Bednar, T.; and Associates, A. 2023. The 2023 Annual Homeless Assessment Report (AHAR) to Congress. Part 1: Point-in-time estimates of homelessness. *The US Department of Housing and Urban Development*.
- Hong, B.; Malik, A.; Lundquist, J.; Bellach, I.; and Kontokosta, C. E. 2018. Applications of machine learning methods to predict readmission and length-of-stay for homeless families: The case of win shelters in new york city. *Journal of Technology in Human Services*, 36(1): 89–104.
- Kuang, D.; Ding, C.; and Park, H. 2012. Symmetric non-negative matrix factorization for graph clustering. In *Proceedings of the 2012 SIAM international conference on data mining*, 106–117. SIAM.
- Kube, A.; 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.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Messier, G.; John, C.; and Malik, A. 2021. Predicting Chronic Homelessness: The Importance of Comparing Algorithms using Client Histories. *Journal of Technology in Human Services*, 1–12.
- Pessach, D.; and Shmueli, E. 2022. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)*, 55(3): 1–44.
- Pokharel, G.; Das, S.; and Fowler, P. 2024. Discretionary trees: understanding street-level bureaucracy via machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22303–22312.
- Rahman, K. S.; and Chelmiss, C. 2022. Learning to Predict Transitions within the Homelessness System from Network Trajectories. In *Proceedings of the 2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 181–189.
- Rahman, K. S.; Zois, D.-S.; and Chelmiss, C. 2023. Bayesian Network Modeling and Prediction of Transitions Within the Homelessness System. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Rodríguez, P.; Bautista, M. A.; Gonzalez, J.; and Escalera, S. 2018. Beyond one-hot encoding: Lower dimensional target embedding. *Image and Vision Computing*, 75: 21–31.
- Sarker, I. H. 2021. Machine learning: Algorithms, real-world applications and research directions. *SN computer science*, 2(3): 160.
- Shinn, M.; Greer, A. L.; Bainbridge, J.; Kwon, J.; and Zuiderveen, S. 2013. Efficient targeting of homelessness prevention services for families. *American journal of public health*, 103(S2): S324–S330.
- Shu, K.; Wang, S.; and Liu, H. 2019. Beyond news contents: The role of social context for fake news detection. In *Proceedings of the twelfth ACM international conference on web search and data mining*, 312–320.
- Tang, J.; and Liu, H. 2012. Unsupervised feature selection for linked social media data. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 904–912.
- Toros, H.; and Flaming, D. 2018. Prioritizing homeless assistance using predictive algorithms: an evidence-based approach. *Cityscape*, 20(1): 117–146.
- United States Department of Housing and Urban Development. 2020. HMIS Data Standards Manual. Retrieved May 26, 2021, from <https://www.hudexchange.info/resource/3824/hmis-data-dictionary/>.
- Vajiac, C.; Frey, A.; Baumann, J.; Smith, A.; Amarasinghe, K.; Lai, A.; Rodolfa, K. T.; and Ghani, R. 2024. Preventing Eviction-Caused Homelessness through ML-Informed Distribution of Rental Assistance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22393–22400.
- VanBerlo, B.; Ross, M. A.; Rivard, J.; and Booker, R. 2021. Interpretable machine learning approaches to prediction of chronic homelessness. *Engineering Applications of Artificial Intelligence*, 102: 104243.
- Wang, S.; Tang, J.; and Liu, H. 2015. Embedded unsupervised feature selection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2019. Fairness constraints: A flexible approach for fair classification. *Journal of Machine Learning Research*, 20(75): 1–42.

Zhang, L.; and Zhang, S. 2017. A unified joint matrix factorization framework for data integration. *arXiv preprint arXiv:1707.08183*.