

Enhancing Equitable Access to AI in Housing and Homelessness System of Care through Federated Learning

Musa Taib*, Jiajun Wu*, Steve Drew, Geoffrey G. Messier

University of Calgary

{musa.taib, jiajun.wu1, steve.drew, gmessier}@ucalgary.ca

Abstract

The top priority of a Housing and Homelessness System of Care (HHSC) is to connect people experiencing homelessness to supportive housing. An HHSC typically consists of many agencies serving the same population. Information technology platforms differ in type and quality between agencies, so their data are usually isolated from one agency to another. Larger agencies may have sufficient data to train and test artificial intelligence (AI) tools but smaller agencies typically do not. To address this gap, we introduce a Federated Learning (FL) approach enabling all agencies to train a predictive model collaboratively without sharing their sensitive data. We demonstrate how FL can be used within an HHSC to provide all agencies equitable access to quality AI and further assist human decision-makers in the allocation of resources within HHSC. This is achieved while preserving the privacy of the people within the data by not sharing identifying information between agencies without their consent. Our experimental results using real-world HHSC data from Calgary, Alberta, demonstrate that our FL approach offers comparable performance with the idealized scenario of training the predictive model with data fully shared and linked between agencies.

Introduction

Emergency housing shelters/agencies are an important part of the housing and homelessness system of care in most North American cities. While originally intended to provide low-barrier access to shelter for short periods, many people now make use of shelters over long periods either continually (*chronic* shelter use) or sporadically (*episodic* shelter use). Most long-term shelter users face multiple physical and mental health challenges (Hopper, Bassuk, and Oliver 2009) that prevent them from making a quick exit from the shelter system without first connecting them with housing and the wrap-around supports necessary to help them address their unique set of challenges. Helping these people exit the shelter quickly is important. Shelters are difficult environments, and the Housing First philosophy recognizes that a person having trouble exiting a shelter should be connected to support as soon as possible (Goering et al. 2011).

Properly matching shelter users to supportive housing is a difficult problem. In most cities, the number of supportive

housing spaces is much less (Coleman 2023) than the number of new people entering homelessness. The vast majority (approximately 85%) of first-time shelter users will exit the shelter quickly and without support (*transitional* shelter use) (Aubry et al. 2013a; Culhane et al. 2007). Therefore, it is the longer-term chronic and episodic shelter users who should be prioritized for the limited supply of supportive housing. However, chronic and episodic shelter use patterns can take months or even years to become evident. As a result, most government definitions of chronic homelessness require a history of homelessness that is one year or longer (Byrne and Culhane 2015). Requiring a person to live in a shelter for this length of time before they can demonstrate a historical record of chronic or episodic homelessness is unacceptable. Living conditions in shelters are difficult and pose further risks to a person’s physical and mental health (Hopper, Bassuk, and Oliver 2009).

Machine learning (ML) has great potential for rapidly identifying which first-time shelter users are at risk of becoming chronic or episodic shelter users in the long term (Toros and Flaming 2018; Purao et al. 2019). Rather than requiring a person to accumulate a record of shelter use stretching over many months or years, ML can identify at-risk people using only the first few months of a person’s shelter record (VanBerlo et al. 2021).

It must be stressed that the purpose of using ML within an HHSC is not to automate the decision of who should be assigned to supportive housing. The established best practice within an HHSC is to support people with a variety of programs, not a single program or automated instrument (Shinn and Richard 2022). Rather than replacing front-line human staff, our tool would be used by that staff to identify possible “under the radar” shelter users who have been missed by other programs. Being identified by our machine learning approach would not automatically qualify someone for housing. Instead, it would initiate a conversation with a support worker to more fully understand that person’s needs.

While ML has great potential, the issue of equitable access to high-quality ML tools within an HHSC is challenging and nuanced. An HHSC typically consists of many not-for-profit agencies of varying sizes. Smaller agencies will usually have fewer IT resources in terms of both expertise and infrastructure. However, IT platform and talent barriers can be overcome by the pooling of resources and/or collab-

*These authors contributed equally.

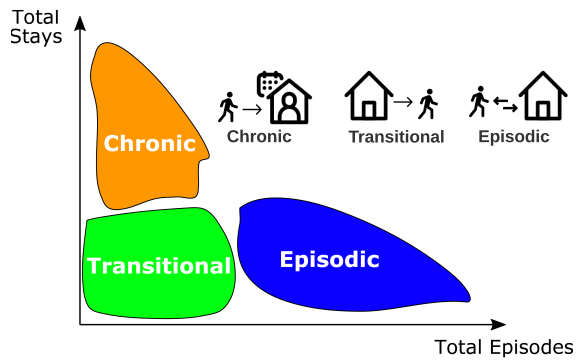


Figure 1: (Total Episode, Total Stay) clusters identifying transitional, chronic and episodic shelter use.

oration between agency IT teams.

The more difficult problem is related to equitable access to ML due to the nature of how data exists within an HHSC. Most HHSCs are made up of a large number of government and not-for-profit community-based services. The same person may interact with multiple agencies, and each agency will record those interactions in its administrative database. Predicting the risk of episodic/chronic homelessness is most accurate when training on data that captures a person's interaction with all the agencies within an HHSC. However, maintaining a single, merged dataset for an entire HHSC is difficult. In the spirit of being a low barrier to entry, many shelters do not require a person to sign data-sharing consent forms. Agency information technology (IT) systems are often incompatible and many municipalities also lack a single system coordinator agency with the expertise and infrastructure to merge data (Matthiesen and Bjørn 2015; Irani et al. 2023). Finally, people are not typically assigned HHSC-wide unique identity numbers which means that merging records between agencies must be done using names and birth dates. Sharing this identifying information across agency boundaries without consent violated the privacy of the people in the data (Murray 2015; Vatsalan, Christen, and Vergyios 2013; Hall and Fienberg 2010).

The fragmentation of data across an HHSC means that agencies typically only have access to the data they collect by themselves. As a result, smaller agencies that see fewer clients will have correspondingly small data sets that make it difficult to effectively train and test an ML model. Equitable access to ML means overcoming this barrier and ensuring that all agencies have access to the same high-quality ML tools, regardless of their data set size. At the same time, the privacy of the people in the data must be respected and identifying information must not be shared between agency boundaries without consent.

Our solution is to utilize a Federated Learning (FL) approach that collectively utilizes the disconnected datasets within the HHSC while respecting the privacy of the people in that data. Our scenario is an example of horizontal partitioning where datasets collected by each individual agency define the partitions. The model will be trained using only shelter stay data (i.e., records of when a person slept in a

shelter). While many other data features (i.e., medical history, demographics, etc.) could potentially be useful predictors, these features are not uniformly collected since not everyone utilizing an HHSC consents to providing information about themselves. There is also a risk of bias if those who do consent to providing additional information become more readily identifiable by the machine learning model. Basing our methods on patterns of shelter use only ensures that our approach can be generalized across the widest possible number of HHSCs and is as equitable as possible.

A FL perspective is also required to label the examples in the historical shelter use dataset used to train and test our model. While many government definitions exist for chronic homelessness (Byrne and Culhane 2015; Messier, Tutty, and John 2022), there are few established definitions for episodic homelessness. Instead, the accepted methodology for labeling a shelter use pattern as chronic, episodic, or transitional is based on the k-means cluster analysis of historical data records (Culhane et al. 2007; Aubry et al. 2013a; Kneebone et al. 2015). The k-means analysis assumes a person's entire record of homelessness is available. This requires a merged dataset of all HHSC records, which is difficult to obtain, as noted above.

The labeling process begins by using the historical shelter access records of a person to determine their number of stays and episodes. An *episode* is defined as a series of shelter stays with gaps of less than 30 days (Byrne and Culhane 2015). As shown in Fig. 1, the k-means algorithm is instructed to create 3 clusters on a scatter plot of the two above values of everyone in the training dataset. The transitional label is assigned to people associated with the low stay/low episode cluster. The episodic labels are assigned to the low stay/high episode and chronic labels are assigned to high stay/low episode clusters, respectively. Our paper will demonstrate how to create Fig. 1 using a federated approach that preserves privacy and does not require merging records.

We consider three scenarios in this paper that achieve different compromises between machine learning model performance and expense:

1. **Centralized:** All agency data has been merged by a central service. The k-means algorithm for assigning the chronic, episodic, and transitional labels to the data is run on this centralized dataset. The machine learning model is trained centrally and then deployed to the individual agencies for their use.
2. **Federated:** Agencies train a single model using a federated approach that avoids the need to centrally merge data. Labels are assigned using a federated k-means algorithm and the model is also trained using a federated approach. The resulting federated model is used by all the individual agencies.
3. **Isolated:** The agencies work in isolation. Each agency runs a k-means labeling algorithm and trains its own machine learning model using only its dataset.

Using HHSC data from Calgary, Alberta, we demonstrate that our FL approach performs almost as well as the much less practical centralized approach. Using per-agency performance metrics, we show that smaller shelters benefit the

most from our collaborative federated approach since they are unable to collect large datasets on their own. Therefore, our work addresses the important data equity issue of ensuring that all shelters have access to well-performing data tools, regardless of their size.

Our contributions are summarized as follows:

- We are the first to demonstrate how FL can be used to provide equitable access to ML tools across an HHSC with fragmented data. This is achieved without sharing data across agency boundaries or violating consent agreements.
- This is the first application of FL to the k-means method of classifying shelter access patterns. While we utilize it in this paper for model training, having these labels is valuable generally for understanding how a population interacts with an HHSC.
- The methods in this paper are demonstrated and verified using real-world historical data from a North American HHSC. While our specific application scenario is predicting chronic and episodic homelessness using the first few months of a person's shelter record, our framework can be generalized to any HHSC administrative data problem.

Related Work

Homelessness Prediction

The Housing First philosophy of rapidly connecting shelter users to housing (Goering et al. 2011; Namian 2020) and the prioritization of people experiencing chronic/episodic homelessness is well established (Aubry 2015; Gaetz et al. 2016). While some alternative methods have been recently proposed (Messier 2023), the most established method for labeling shelter access patterns as chronic, episodic or transitional is using a k-means clustering analysis of historical data records (Byrne and Culhane 2015; Aubry et al. 2013b; Kneebone et al. 2015). To date, much of the machine learning work related to homelessness has been concentrated on predicting a first entrance or re-entrance to homelessness (Clark et al. 2021; DiGuseppi et al. 2020; Chan et al. 2017; Kube, Das, and Fowler 2019; Gao, Das, and Fowler 2017; Hong et al. 2018; Shinn et al. 2013; Byrne, Montgomery, and Fargo 2019). The specific prediction of chronic homelessness has also been explored (VanBerlo et al. 2021; Toros and Flaming 2018; Purao et al. 2019). However, all of these studies assume a single dataset. None have explored how to overcome the practical issues that arise when working with data that has been fragmented across an HHSC.

Federated Learning

Federated Learning (FL) is a decentralized and privacy-preserving machine learning approach that involves training local models in each agency/device using local data, then aggregating the local models using only the model weights (Yang et al. 2019; Li et al. 2020; Lim et al. 2020). First introduced by McMahan et al. (2017), FL enables collaborative machine learning without the need for raw data to be shared with a central authority. In recent years, FL has seen

a surge in the healthcare sector where patient data privacy is a paramount concern (Rieke et al. 2020; Brisimi et al. 2018; Xu et al. 2021; Sheller et al. 2020).

FL for Risk Prediction Healthcare, a field highly analogous to the services provided by an HHSC, has adopted FL widely for building collaborative prediction models with administrative data (Brisimi et al. 2018; Dayan et al. 2021), in particular, Electronic Health Records (EHRs). For instance, FL has been used to train prediction models for cardiovascular conditions (Yaqoob et al. 2023), clinical outcomes of COVID-19 patients (Dayan et al. 2021), as well as multi-party diabetes mellitus risk prediction (Su et al. 2023). Despite FL's capability to securely handle sensitive data, no existing studies are using FL with shelter data for risk prediction in the field of Homelessness. This study aims to close this gap by introducing a practical FL training framework for shelters to collaboratively predict the risk of episodic and chronic homelessness without sharing the data.

Dataset

This study utilizes anonymized shelter data provided by the Calgary Homeless Foundation (CHF). The protocol governing the anonymization, secure storage, and analysis of this secondary dataset was approved by the University of Calgary Conjoint Ethics Review Board (REB-19-0095). The CHF is an organization that provides and supports a data record system utilized by most of the housing and homelessness service agencies in Calgary, Canada. This allows the CHF to aggregate and merge records of service use for the same group of people accessing services across the system of care using privacy-preserving methods (Vatsalan, Christen, and Verykios 2013). As a result, the anonymized dataset used for this paper represents the *Centralized* scenario described in the Introduction Section. The dataset stretches from January 1, 2009 to December 31, 2019 and contains timestamped records of when a person using the HHSC slept at an emergency shelter. There are 6,840,069 sleep records for 50,455 individuals accessing 8 different shelters.

Restricting the machine learning model to predict episodic, chronic, and transitional shelter access patterns using sleep event data only is challenging. There are many other data features sometimes recorded by HHSC services that may be useful in a predictive model (i.e., demographic information, medical history, employment status, family relationships, etc.). However, it is important to stress that this information is not always collected. Demographic information may only be recorded when accessing certain programs. Low barrier-to-entry emergency shelters may not require clients to disclose any information to access sleep services on a cold night. In contrast, a record of sleeping in a shelter is consistently and uniformly collected for all HHSC users. Focusing on this data feature ensures all people are represented equitably in the data. While specific data recording practices may vary between HHSCs, shelter sleep data is recorded by most, which means the methods presented in this paper can be easily adopted in many jurisdictions.

Data Preprocessing

The dataset utilized for this study has a series of time-stamped records indicating when a person slept in an emergency shelter. This data is pre-processed as follows:

1. The sleep data records for each client are truncated after T_o days. This is the observation window used to predict whether the client will become chronic/episodic.
2. The T_o day observation window is then subdivided into T_b equal segments known as time bins, where the number of sleeps in each bin is summed. For example, a client with a $T_o = 120$ -day observation window divided into $T_b = 10$ time bins would have each bin containing the sum of their shelter stays over a period of 12 days. Each of these time bin values is a different input data feature for the machine learning algorithm.
3. As mentioned before our data is a series of timestamps however that does not provide much useful information to ML models. Therefore, we engineered two features to reflect the number of sleep and episodes during the observation window to increase the total number of features. We counted the number of days that a client slept at a shelter over a specified period, and an episode is defined as a period of 30 or more days between consecutive instances of sleeping.

Note that the HHSC dataset was pre-screened by the CHF, so it was unnecessary to add preprocessing to deal with invalid data entries. By default, the data from the CHF is linked across all HHSC agencies and corresponds to the Centralized dataset scenario. To evaluate the Federated and Isolated scenarios, data is unlinked so that a person is assigned a new client identification (ID) number each time they access a new agency.

Labeling

The labeling process begins by using the historical shelter access records of an individual. A person’s pattern of shelter access is labeled as transitional, chronic, or episodic by first creating the 2-tuple (N_S, N_E) for that person where N_S is that person’s total number of shelter stays, and N_E is the person’s total number of shelter episodes. N_S and N_E are calculated using the person’s first T_p days of HHSC interaction.

In a machine learning context, T_p corresponds to the *prediction window* since this is the window over which a person’s specific pattern of shelter access will manifest. In previous studies (Culhane et al. 2007; Aubry et al. 2013a; Kneebone et al. 2015), T_p has been equal to the total number of days of data available for each person. However, in the CHF data described in the Data Section, this can be several years, which is an overly ambitious prediction window given that we would like to restrict the observation window T_o to only a few months. We will demonstrate in the Hyperparameter Selection Section that reducing T_p to be on the order of 1-2 years improves performance while still providing a very impressive prediction horizon.

Once the 2-tuples are determined for each person in the data, the k-means algorithm with 3 clusters is used to asso-

Algorithm 1: DecentralizedLabeling

Input: \mathcal{K} is the list of agencies; $k \in \mathcal{K}$ is an agency; $|\mathcal{D}_k|$ is the size of the dataset for agency k . $\mu_c^k, \mu_e^k, \mu_t^k$ are chronic, episodic, and transitional cluster centroids on agency k ; μ_c, μ_e, μ_t are global chronic, episodic, and transitional cluster centroids used for labeling.

Output: μ_c, μ_e, μ_t .

- 1: **for** each agency $k \in \mathcal{K}$ **do**
 - 2: Initialize $\mu_c^k, \mu_e^k, \mu_t^k$ from \mathcal{D}_k
 - 3: Apply k-means clustering for \mathcal{D}_k
 - 4: Update $\mu_c^k, \mu_e^k, \mu_t^k$
 - 5: **end for**
 - 6: $\mu_c, \mu_e, \mu_t \leftarrow \sum_{k=1}^{\mathcal{K}} \frac{|\mathcal{D}_k|}{\sum_{i=1}^{\mathcal{K}} |\mathcal{D}_i|} \mu_c^k, \mu_e^k, \mu_t^k$
 - 7: **return** μ_c, μ_e, μ_t
-

ciate each person with a centroid that corresponds to transitional, chronic, and episodic shelter use as shown in Fig. 1. For the Centralized scenario, the 2-tuples are created for everyone in the data using the fully merged CHF dataset. For the Isolated case, a person’s interaction with multiple agencies is unlinked and treated as different individuals as described in the Data Preprocessing Section. Each agency performs its own k-means labeling process for the unlinked individuals in their own datasets.

For the Federated scenario, we propose a new labeling approach, Decentralized k-means, as shown in Algorithm 1, that allows agencies to benefit from pooling their data without the technical complexity and privacy concerns of fully linking their datasets. Each agency performs k-means clustering on its own datasets for the Isolated case. The clusters are then forwarded to a central location that performs a weighted average. These averaged clusters are then shared back with the agencies who can use them to re-classify their 2-tuples.

Methodology

Problem Formulation

We formulate the allocation of the homelessness support problem as a multi-class prediction task. Consider a dataset \mathcal{D} , which is a $m \times n$ matrix where m is the number of clients and n is the number of features. The sleep pattern of an individual client is split into T_b time bins in addition to two engineered features: total number of sleeps and total number of episodes. Therefore, each client would have an input feature vector \mathbf{f} with length $n = T_b + 2$. Our task is to train a classifier $C(\cdot)$, which produces the prediction results for a given \mathbf{f} ; $\hat{y} = C(\mathbf{f})$, where $\hat{y} \in [\text{chronic, episodic, transitional}]$ refers to the output labels.

Training Details

Fig. 2 shows the training scenarios of the three models. Training for Centralized and Isolated scenarios involves the preprocessing steps described in the Data Preprocessing Section, dataset labeling as described in the Labeling Section, and updating the machine learning model as shown in




Figure 2: Illustration of the three scenarios for training predictive models in the paper: Federated, Centralized, and Isolated.

Algorithm 2: AgencyTraining. The Local Prediction Model Training for each agency.

Input: w is the machine learning model; η is the learning rate; \mathcal{E} is the number of training epochs; B is the training batch size; \mathcal{L} is the loss function of training.

Output: The trained agency model.

- 1: $\mathcal{B} \leftarrow \text{split } \mathcal{D}_k \text{ into batches of size } B$
- 2: **for** each local epoch i from 1 to \mathcal{E} **do**
- 3: **for** each batch $b \in \mathcal{B}$ **do**
- 4: $w \leftarrow w - \eta \nabla \mathcal{L}(w; b)$
- 5: **end for**
- 6: **end for**
- 7: **return** w

Algorithm 2. The only difference is that Centralized and Isolated scenarios utilize linked and unlinked data described in the Data Preprocessing Section, respectively.

Our proposed FL framework is based on the FedAvg algorithm, following the steps below:

- **Step 1:** Apply Decentralized k-means from Algorithm 1 to generate the labels for the training data.
- **Step 2:** The central server broadcasts the global model to \mathcal{K} available local agencies.
- **Step 3:** Each agency updates and uploads its local model to the server after local training.
- **Step 4:** The server aggregates all uploaded local models into the new global model for the next round. Repeat **Steps 2-4** until the global model converges.

The pseudocode for the FL scenario is shown in Algorithm 3.

Evaluation

We explain the experimental setup of our 3 proposed methods: Centralized, Isolated, and Federated.

Algorithm 3: Federated Prediction Model Training

Input: \mathcal{K} is the list of agencies; \mathcal{D} is the list of dataset for all agencies; $|\mathcal{D}_k|$ is the size of the dataset for agency k ; \mathcal{T} is the number of communication rounds; w_t^k is the global model for agency k at round t ; B is the batch size; θ is the clusters.

Output: The trained global model parameters.

- 1: DecentralizedLabeling(\mathcal{K}, \mathcal{D}) ▷ Alg. 1
- 2: Server initialize w_0 ;
- 3: **Training Starts:**
- 4: **for** each round t from 1 to \mathcal{T} **do**
- 5: **for** each agency $k \in \mathcal{K}$ in parallel **do**
- 6: $w_t^k \leftarrow \text{AgencyTraining}(k, \mathcal{D}_k, w_{t-1})$ ▷ Alg. 2
- 7: **end for**
- 8: $w_{t+1} \leftarrow \sum_{k=1}^{\mathcal{K}} \frac{|\mathcal{D}_k|}{\sum_{i=1}^{\mathcal{K}} |\mathcal{D}_i|} w_t^k$
- 9: **end for**

Experimental Setup

The experiments were conducted in a simulated setup on an Apple M1 Pro MacBook with 16GB of RAM¹. The neural networks trained at each simulated agency were isolated to ensure privacy for isolated and federated models. There was no direct inter-agency communication except for controlled weight sharing between training rounds in federated cases. There was no communication for isolated cases.

Train/Test Splits

We apply a normalization process to the dataset using z-scoring (Pedregosa et al. 2011). This results in a standardized score for each data point, reflecting how many standard deviations it is from the column average.

Following the normalization, the dataset is divided into stratified training and testing sets, adhering to an 80/20 split. This strategy guarantees that both sets maintain proportional representations of clients across different agencies and la-

¹The details of the model and code are available at <https://github.com/MusaJTaib/FederatedLearning>

bels. For the Federated and Isolated models, the test set labels are synchronized with those from the centralized model.

Machine Learning Models

Multi-layer Perceptrons (MLPs) are a class of neural networks characterized by their stacked arrangement of fully connected layers, which makes them particularly adept at extracting and learning complex patterns from the input data (Goodfellow, Bengio, and Courville 2016). The MLP model used for this paper features three ReLU-activated hidden layers (Glorot, Bordes, and Bengio 2011) of sizes 512, 128, and 16, forming a structure capable of identifying and processing complex patterns within the dataset. This configuration, with its hierarchical design, is adept at capturing and learning the intricate relationships present in the input data. To facilitate model generalization and curtail overfitting, dropout layers with rates ranging from 0.1 to 0.4 are judiciously applied across the dense layers (Srivastava et al. 2014). In the final output layer, a sigmoid function (Goodfellow, Bengio, and Courville 2016) is employed to handle multi-class classification.

The training for each model is performed as described in Algorithm 2. The models employed the Adam optimizer (Kingma and Ba 2014) with a learning rate $\eta = 0.02$. For loss calculation, categorical cross-entropy is selected, aligning with the multi-class nature of the output. Optimization is performed via mini-batch gradient descent with a batch size $B = 500$ and the models are trained for a total of 200 epochs (\mathcal{E}).

The performance of the model can vary as a result of random weight initializations. Therefore, each MLP model was trained 10 times on different initializations for all three cases (federated, isolated, and centralized). The results for all 10 rounds were then averaged to get a more generalized picture of the results.

The aforementioned details apply uniformly across all three models, with the federated model incorporating additional parameters to tailor the training process. According to Algorithm 3, the federated model undergoes $\mathcal{T} = 75$ communication rounds. With $\mathcal{K} = 8$ each representing a unique agency contributing to the model. Furthermore, each client undergoes $\mathcal{E} = 15$ training epochs on their local dataset for every round.

Hyperparameter Selection

The Data Preprocessing and Labeling Sections present some important hyperparameters that have a significant impact on model performance. Increasing the size of the observation window, T_o , will generally improve model performance. However, a long observation window is at odds with our objective of identifying at-risk HHSC users as soon as possible. Therefore, the smallest possible T_o that yields acceptable performance must be selected.

Increasing the number of time bins, T_b , improves the model’s ability to capture temporal trends and provide a better fit to the data. In general, a small value of T_b will underfit and a large value of T_b will overfit the data (Taib and Messier 2023).

T_b (Bins)	T_o (Days)	T_p (Days)	Precision	Recall
5	548	120	0.618	0.735
		90	0.568	0.691
		60	0.529	0.648
	730	120	0.584	0.696
		90	0.547	0.654
		60	0.504	0.609
	913	120	0.549	0.666
		90	0.514	0.631
		60	0.479	0.593
10	548	120	0.622	0.729
		90	0.571	0.683
		60	0.530	0.634
	730	120	0.608	0.686
		90	0.545	0.656
		60	0.499	0.604
	913	120	0.560	0.658
		90	0.511	0.632
		60	0.481	0.593

Table 1: Precision and Recall values for Centralized model with different time bins (T_b), observations (T_o) and prediction (T_p) windows.

Model	Linked Data	Precision	Recall
Centralized	✓	0.5713	0.6826
Federated		0.5379	0.6495
Isolated		0.5061	0.4656

Table 2: Prediction performance for the Centralized, Federated, and Isolated scenarios.

Finally, model performance will improve with a decreasing prediction window T_p since a smaller T_p improves the relationship between the labels and the pattern of shelter use early in a person’s HHSC record. However, T_p should still be made as large as possible to allow the longer term patterns of chronic and episodic homelessness to develop.

Table 1 shows model performance for the Centralized scenario for various values of T_o , T_b and T_p . As noted above, the best performance is achieved with the largest values of T_o and T_b and the smallest value of T_p . In the remainder of the paper, the values of $T_b = 10$ and $T_p = 548$ days will be used. However, a value of $T_o = 90$ days will be used instead of 120 days to achieve a compromise between performance and rapid identification.

Comparing Centralized, Federated and Isolated Prediction

Using the hyper-parameters from the Hyperparameter Selection Section, Table 2 shows precision and recall for all three data scenarios averaged across 10 randomized initializations.

As expected, the Centralized scenario achieves the best

Model	Transitional	Episodic	Chronic
C-P	0.90	0.43	0.30
F-P	0.92	0.27	0.41
I-P	0.80	0.28	0.42
C-R	0.82	0.43	0.72
F-R	0.68	0.59	0.69
I-R	0.74	0.45	0.20

Table 3: Prediction performance per class for the Centralized, Federated, and Isolated scenarios. C-, F-, and I- denote Centralized, Federated, and Isolated respectively. P: precision, R: recall.

performance in Table 2 due to the model being able to utilize a fully linked dataset. However, the Federated model achieves performance that comes within 4% of the Centralized model. This is excellent performance, especially since the federated methods avoid the privacy and technical barriers of linking data.

Looking at Table 3 we can also see another advantage of using FL over the current local shelter implementation. The per class precision and recall for the FL case are in general much better than for local classes and fall in line with previous trends. This is quite important for this case because of the highly imbalanced nature of the dataset. As overall higher results shown before could have simply been attributed to a single class’s performance improvement. However, it is quite obvious that clients in all three classes are seeing the benefits of this improvement making the model even more fairer to use compared to the current local implementations.

Performance achieved by the Isolated scenario is far below Federated and Centralized, especially in terms of recall. This shortfall can be explained by the per-agency Federated and Isolated scenario performance results presented in Table 4. Each row in this table includes the agency ID number, the number of individuals in that agency’s dataset, and the precision and recall achieved by the Federated and Isolated scenarios.

Table 4 demonstrates that it is the smaller shelters (shelters with capacity below 12,000) that perform particularly poorly in the Isolated case due to the very small size of their datasets. An exception to this trend is noted in the case of agency 213. Given its limited size (293 clients) and the consequently small test sample (approximately 70 clients), the results for this agency are deemed insufficient for a reliable evaluation of model performance and are considered an outlier.

Figure 3 compares the Federated and Isolated cases using the F1 score. In this figure, it is evident that agencies with large datasets (notably 330, 188, and 55) can achieve respectable performance using their own datasets in isolation. This further underscores that the Federated approach has the largest impact on shelters with smaller operations.

Agency	Clients	F-P	F-R	I-P	I-R
4	2386	0.411	0.641	0.203	0.397
13	7063	0.443	0.659	0.343	0.388
55	12017	0.502	0.67	0.57	0.411
188	15065	0.605	0.549	0.499	0.452
213	293	0.381	0.508	0.38	0.636
225	1207	0.515	0.536	0.249	0.386
330	28523	0.581	0.716	0.629	0.568
333	11713	0.473	0.607	0.341	0.353

Table 4: Experiments on the real-world dataset. F- and I- denote Federated and Isolated, respectively. P: precision, R: recall.




Figure 3: F1 Scores for Federated and Isolated cases

Discussion

Before discussing the results in the Evaluation Section, it is important to provide context regarding how a machine learning tool should be utilized within an HHSC. As noted in the Introduction Section, machine learning should never be the only method for connecting people to housing. Instead, it is used as part of a range of programs for engaging with HHSC users (Shinn and Richard 2022), including “diversion” programs specifically designed for making contact with new HHSC entrants.

While improving the absolute performance of the models presented in the Evaluation Section is always desirable, these performance levels are acceptable if the models are used in the context of supporting a robust diversion program with human front-line staff. Performance could also be further improved if the methods in this paper are applied in an HHSC where other data features, such as demographic information, are collected.

Centralized vs. Federated The results from Table 2 show that the centralized model outperforms the other two models. The Centralized scenario is ideal and may be achievable in some HHSCs that have a centralized agency capable of collecting data, as is the case in Calgary. However, most HHSCs are better represented by the Isolated scenario where agen-

cies would need to operate independently without collaboration. The Federated approach, as proposed in this paper, is posited as a viable compromise, offering a balance between prediction performance and minimizing the amount of centralized IT infrastructure required. It can also be implemented without requiring the people using the HHSC to provide consent to share data across agency boundaries.

Achieving Equity with Federated Learning One significant challenge within the HHSC network is the disparity in data analytics capabilities between large and small agencies. Smaller agencies often struggle with insufficient datasets. This limitation hampers their ability to identify at-risk clients accurately and tailor their services effectively. Looking at the results in Figure 3, it is quite clear that our federated model has a substantial gain in performance over the isolated case. A breakdown of the per agency performance keeping in mind the number of clients per agency in table 4 shows exactly where the federated model gains its performance. Notably, smaller shelters, those with fewer than 10,000 clients, which struggle to adequately train their models due to insufficient data, benefit markedly from the aggregated insights available through larger shelters. Conversely, larger shelters observed minimal performance gains, as their pre-existing data volumes were already sufficient to effectively train their models.

The federated model presents a solution to this problem by enabling smaller agencies to leverage larger, aggregated models trained across multiple data sources without directly sharing sensitive information. FL allows for the collective improvement of model performance, benefiting all participants in the network. When using machine learning for social good, equity is important. This means that all people accessing an HHSC should benefit from the same high-quality tools regardless of whether they choose to interact with a large or small agency. The fact that our proposed FL approach benefits smaller agencies, in particular, shows it is an important tool for achieving this equity goal.

Federated models reduce the need for data linkage The conventional approach to enhancing machine learning models across different agencies involves the linkage of data to create a centralized dataset. This process not only raises significant privacy concerns but also involves logistical complexities and substantial resources dedicated to ensuring the secure and accurate merging of sensitive information. The federated model offers a compelling advantage by eliminating the necessity for direct data linkage. By training models on decentralized data and aggregating the learned models, FL circumvents the challenges associated with data consolidation. This approach significantly reduces the complexity and cost associated with data linkage, reallocating those resources toward improving service delivery and operational efficiencies. Moreover, FL upholds the privacy and confidentiality of individual records, addressing one of the critical barriers to data sharing among agencies.

Conclusion

In this work, a critical gap in current homelessness management strategies is addressed by using a machine learning

approach that respects the privacy and autonomy of individual agencies while leveraging data from across the HHSC. In jurisdictions where raw data sharing across agencies is challenged by privacy concerns or logistical challenges, our proposed FL model demonstrated virtually the same performance as that of a centralized data approach. Compared to the status quo in the domain, our federated approach greatly enhances the ability of smaller shelters to make predictions and also promotes equity across the HHSC. By demonstrating the feasibility and effectiveness of federated learning in this context, this study lays the groundwork for future research and development of ML applications within the social services sector.

Acknowledgements

The authors would like to acknowledge the support of Making the Shift, the Calgary Homeless Foundation, and the Government of Alberta. This study is based in part on data provided by Alberta Seniors, Community and Social Services. The interpretation and conclusions contained herein are those of the researchers and do not necessarily represent the views of the Government of Alberta. Neither the Government of Alberta nor Alberta Seniors, Community, and Social Services express any opinion related to this study. This research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant RGPIN-2024-03954.

References

- Aubry, T. 2015. *Screening for housing first*. Canadian Observatory on Homelessness.
- Aubry, T.; Farrell, S.; Hwang, S. W.; and Calhoun, M. 2013a. Identifying the Patterns of Emergency Shelter Stays of Single Individuals in Canadian Cities of Different Sizes. *Housing Studies*, 28(6): 910–927. Publisher: Routledge. eprint: <https://doi.org/10.1080/02673037.2013.773585>.
- Aubry, T.; Farrell, S.; Hwang, S. W.; and Calhoun, M. 2013b. Identifying the Patterns of Emergency Shelter Stays of Single Individuals in Canadian Cities of Different Sizes. *Housing Studies*, 28(6): 910–927.
- Brisimi, T. S.; Chen, R.; Mela, T.; Olshevsky, A.; Paschalidis, I. C.; and Shi, W. 2018. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics*, 112: 59–67.
- Byrne, T.; and Culhane, D. P. 2015. Testing Alternative Definitions of Chronic Homelessness. *Psychiatric Services*, 66(9): 996–999. Publisher: American Psychiatric Publishing.
- Byrne, T.; Montgomery, A. E.; and Fargo, J. D. 2019. Predictive modeling of housing instability and homelessness in the Veterans Health Administration. *Health Services Research*, 54(1): 75–85. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1475-6773.13050>.
- Chan, H.; Rice, E.; Vayanos, P.; Tambe, M.; and Morton, M. 2017. Evidence From the Past: AI Decision Aids to Improve

- Housing Systems for Homeless Youth. *AAAI Fall Symposium Series*, 149–157.
- Clark, A. Y.; Blumenfeld, N.; Lal, E.; Darbari, S.; Northwood, S.; and Wadpey, A. 2021. Using K-Means Cluster Analysis and Decision Trees to Highlight Significant Factors Leading to Homelessness. *Mathematics*, 9(2045): 2045. Publisher: MDPI AG.
- Coleman, A. 2023. Diversion: An Overview of Emergency Shelter Diversion as a Practice and the Local Context in Waterloo Region. Available online: URL (Accessed on Date).
- Culhane, D. P.; Metraux, S.; Park, J. M.; Schretzman, M.; and Valente, J. 2007. Testing a typology of family homelessness based on patterns of public shelter utilization in four U.S. jurisdictions: Implications for policy and program planning. *Housing Policy Debate*, 18(1): 1–28. Publisher: Routledge [eprint: https://doi.org/10.1080/10511482.2007.9521591](https://doi.org/10.1080/10511482.2007.9521591).
- Dayan, I.; Roth, H. R.; Zhong, A.; Harouni, A.; Gentili, A.; Abidin, A. Z.; Liu, A.; Costa, A. B.; Wood, B. J.; Tsai, C.-S.; et al. 2021. Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature medicine*, 27(10): 1735–1743.
- DiGuseppi, G. T.; Davis, J. P.; Leightley, D.; and Rice, E. 2020. Predictors of Adolescents' First Episode of Homelessness Following Substance Use Treatment. *Journal of Adolescent Health*, 66(4): 408–415.
- Gaetz, S.; Dej, E.; Richter, T.; and Redman, M. 2016. The State of Homelessness in Canada 2016. *Canadian Observatory on Homelessness Press*.
- Gao, Y.; Das, S.; and Fowler, P. J. 2017. Homelessness Service Provision: a Data Science Perspective. *The AAAI-17 Workshop on AI and Operations Research for Social Good*, 20–24.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep Sparse Rectifier Neural Networks. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*.
- Goering, P. N.; Streiner, D. L.; Adair, C.; Aubry, T.; Barker, J.; Distasio, J.; Hwang, S. W.; Komaroff, J.; Latimer, E.; Somers, J.; and Zabkiewicz, D. M. 2011. The At Home/Chez Soi trial protocol: a pragmatic, multi-site, randomised controlled trial of a Housing First intervention for homeless individuals with mental illness in five Canadian cities. *BMJ Open*, 1(2): e000323. Publisher: British Medical Journal Publishing Group Section: Health services research.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Hall, R.; and Fienberg, S. E. 2010. Privacy-Preserving Record Linkage. In Domingo-Ferrer, J.; and Magkos, E., eds., *Privacy in Statistical Databases*, Lecture Notes in Computer Science, 269–283. Berlin, Heidelberg: Springer. ISBN 978-3-642-15838-4.
- 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. Publisher: Routledge [eprint: https://doi.org/10.1080/15228835.2017.1418703](https://doi.org/10.1080/15228835.2017.1418703).
- Hopper, E.; Bassuk, E.; and Oliver, J. 2009. Shelter from the Storm: Trauma-Informed Care in Homelessness Services Settings. 80–100. Accessed December 2012.
- Irani, Z.; Abril, R. M.; Weerakkody, V.; Omar, A.; and Sivarajah, U. 2023. The impact of legacy systems on digital transformation in European public administration: Lesson learned from a multi case analysis. *Government Information Quarterly*, 40(1): 101784.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- Kneebone, R. D.; Bell, M.; Jackson, N.; and Jadidzadeh, A. 2015. Who Are the Homeless? Numbers, Trends and Characteristics of Those Without Homes in Calgary. *The School of Public Policy Publications (SPPP)*, 8. Place: Calgary, Canada Publisher: University of Calgary, The School of Public Policy Section: Research Papers.
- Kube, A.; Das, S.; and Fowler, P. J. 2019. Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 622–629. Number: 01.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3): 50–60.
- Lim, W. Y. B.; Luong, N. C.; Hoang, D. T.; Jiao, Y.; Liang, Y.-C.; Yang, Q.; Niyato, D.; and Miao, C. 2020. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3): 2031–2063.
- Matthiesen, S.; and Bjørn, P. 2015. Why replacing legacy systems is so hard in global software development: An information infrastructure perspective. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, 876–890.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Messier, G.; Tutty, L.; and John, C. 2022. The Best Thresholds for Rapid Identification of Episodic and Chronic Homeless Shelter Use. *International Journal on Homelessness*, 2(1): 37–47. Number: 1.
- Messier, G. G. 2023. A Simpler Method for Understanding Emergency Shelter Access Patterns.
- Murray, J. S. 2015. Probabilistic Record Linkage and Deduplication after Indexing, Blocking, and Filtering. *Journal of Privacy and Confidentiality*, 7(1). Number: 1.
- Namian, D. 2020. Governing homelessness through instruments: a critical perspective on housing first's policy instrumentation. *Critical Policy Studies*, 14(3): 303–318. Publisher: Routledge [eprint: https://doi.org/10.1080/19460171.2019.1598881](https://doi.org/10.1080/19460171.2019.1598881).

- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.
- Purao, S.; Garfield, M.; Gu, X.; and Bhetwal, P. 2019. Predicting the Slide to Long-Term Homelessness: Model and Validation. In *2019 IEEE 21st Conference on Business Informatics (CBI)*, volume 01, 31–40. ISSN: 2378-1971.
- Rieke, N.; Hancox, J.; Li, W.; Milletari, F.; Roth, H. R.; Albarqouni, S.; Bakas, S.; Galtier, M. N.; Landman, B. A.; Maier-Hein, K.; et al. 2020. The future of digital health with federated learning. *NPJ digital medicine*, 3(1): 119.
- Sheller, M. J.; Edwards, B.; Reina, G. A.; Martin, J.; Pati, S.; Kotrotsou, A.; Milchenko, M.; Xu, W.; Marcus, D.; Colen, R. R.; et al. 2020. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*, 10(1): 12598.
- 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. Publisher: American Public Health Association.
- Shinn, M.; and Richard, M. K. 2022. Allocating Homeless Services After the Withdrawal of the Vulnerability Index—Service Prioritization Decision Assistance Tool. *American Journal of Public Health*, 112(3): 378–382. Publisher: American Public Health Association.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15: 1929–1958.
- Su, Y.; Huang, C.; Zhu, W.; Lyu, X.; and Ji, F. 2023. Multi-party diabetes mellitus risk prediction based on secure federated learning. *Biomedical Signal Processing and Control*, 85: 104881.
- Taib, M.; and Messier, G. G. 2023. Efficient Observation Time Window Segmentation for Administrative Data Machine Learning.
- Toros, H.; and Flaming, D. 2018. Prioritizing Homeless Assistance Using Predictive Algorithms: An Evidence-Based Approach. *Cityscape*, 20(1): 117–146. Publisher: US Department of Housing and Urban Development.
- VanBerlo, B.; Ross, M. A. S.; Rivard, J.; and Booker, R. 2021. Interpretable machine learning approaches to prediction of chronic homelessness. *Engineering Applications of Artificial Intelligence*, 102: 104243.
- Vatsalan, D.; Christen, P.; and Verykios, V. S. 2013. A taxonomy of privacy-preserving record linkage techniques. *Information Systems*, 38(6): 946–969.
- Xu, J.; Glicksberg, B. S.; Su, C.; Walker, P.; Bian, J.; and Wang, F. 2021. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5: 1–19.
- Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2): 1–19.
- Yaqoob, M. M.; Nazir, M.; Khan, M. A.; Qureshi, S.; and Al-Rasheed, A. 2023. Hybrid classifier-based federated learning in health service providers for cardiovascular disease prediction. *Applied Sciences*, 13(3): 1911.