Project HOPE: Homelessness Outreach Planning Effort

Noah Zweben, Elena Horton, and Neeraj Chandra
MS/MBA Program, Harvard University, CS 288

Abstract

This paper presents Project HOPE, an online AI-enabled tool to assist outreach workers in their care of rough sleepers (unhoused individuals living on the streets). Using a block-level predictor of rough sleeper counts, HOPE optimally routes outreach workers over the course of their shift, maximizing both the total number of rough sleepers treated and the geographic area covered. We compare HOPE to various naive and heuristic approaches and find that HOPE performs strongly – warranting further study in the intersection of outreach work and AI.

Keywords: AI4SG; homelessness outreach; optimization; traveling salesman routing

1 Context and motivation

Rough sleeping, the most visible form of homelessness, refers to individuals sleeping in places that are not designated for living including the street, tents, and cars. Rough sleepers are exposed to crime, unsanitary conditions, and weather, posing significant risks to health and welfare. *Outreach coordinators* are the front-line mechanism for caring for rough sleepers. These workers, sometimes volunteers and sometimes staffed by local non-profits, canvas cities by foot, bike, or car to find rough sleepers and engage with them, earning trust and informing them of helpful resources, such as housing, medical, and employment services. It often takes dozens of conversations to build enough trust to enable a meaningful support relationship, and individual interventions can involve lengthy conversations. Outreach workers must search over a large geographic area with often limited employees and engage in time-intensive care. Currently, outreach workers use past experience and intuition to direct their surveys.

1.1 Problem

Outreach workers are currently plagued by three persistent problems. First, non-profits coordinating outreach programs are often very under-resourced. For example, the city of San Antonio's homeless outreach program is currently at 50% capacity. A lack of personnel impacts the ability to canvas the entire city, so workers choose to take routes based on simplistic planning or intuition. Secondly, with a haphazardly planned approach, outreach workers exhibit very little path variation day to day. They often take repeated routes, which limits their ability to explore new areas and encounter disparate groups of rough sleepers. This affects the diversity of rough sleepers encountered, as disparate vulnerable populations – such as youth or women – may reside in different areas. In addition to being sub-optimal with regards to diversity, it is likely these paths are also sub-optimal in terms of the total number of rough sleepers encountered.

These issues were discovered through interviews with current outreach workers in the United States. To address these problems and assist outreach workers in discovering rough sleepers, we propose Project HOPE (Homelessness Outreach Planning Effort). We begin this project by asking the question– *can we use AI to maximize the number of rough sleepers an outreach worker can engage with on their daily shift?*

2 Related Work

A few academics have applied data science techniques to the problem of homelessness, although, to our knowledge, none have specifically tried to improve the effectiveness of outreach workers. *VanBerlo et al.* applied a novel model, *HIFIS-RNN-MLP*, to predict chronic homelessness likelihood based on client data from Canadian homeless shelters. *Shah et al.* use both LSTM and ARIMA models to predict likelihood of victimizations for specific individuals experiencing homelessness based on their age, sex, race, and location using combined municipal datasets within Los Angeles. The most similar is *Wilde et al.*, which trains a classifier to prioritize crowd-sourced reports of rough sleepers in need of assistance, versus predicting where to go in the absence of these reports. We take these examples as helpful context for data usage and complementary to our work.

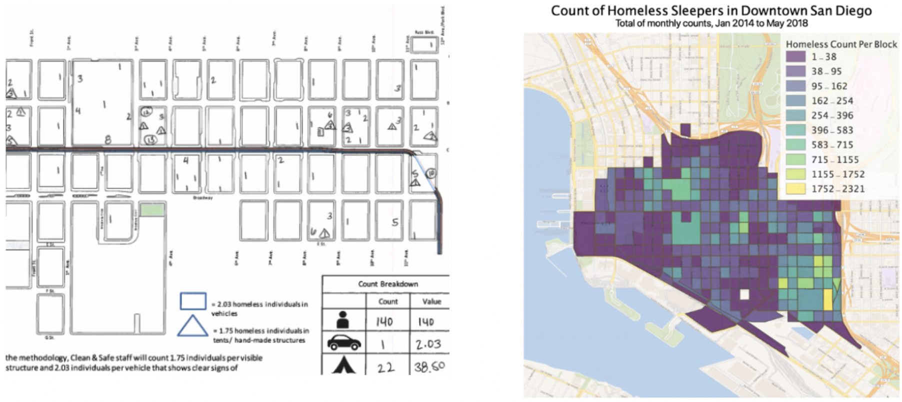

Figure 1: An example of a data entry by the San Diego Downtown Partnership. These hand-marked maps were digitized by the Civic Knowledge Agency. The right figure is the aggregated counts by block provided by the dataset authors.

Similar work on predicition-optimized routing has been undertaken in the policing domain, such as by *Guevara et. al.* These studies were sources of inspiration. However, some key differences between our access to ground-truth data and biased crime data are covered in the **Data section**.

3 Data

We obtain our data from the San Diego Downtown Partnership (1). Between the years of 2014-2018, the Partnership coordinated monthly sweeps of a 287 block area of downtown San Diego (8). On each sweep, volunteers hand-recorded the number of rough sleepers encountered on every block (see *Figure 1* for an example of one of these entries). This data has since been digitized by the *Civic Knowledge Agency*, allowing us to work with these 88 months of data organized in a time series for each of the 287 blocks. Additional data from 2018-2022 has yet to be digitized and is excluded from our analysis. These surveys are unique in that they **provide full coverage of the downtown area**. This enables us to treat the count data as ground truth, largely free from the location-selection bias of the surveyors. This stands in stark contrast to policing and crime data, in which much of the data is predicated on prior biased decisions, like on whom and where surveillance and measurement efforts have been concentrated.

A natural question we had was how well exogenous factors – for example, weather, pedestrian traffic, the prevalence of businesses, etc. – could be correlated with rough sleeper counts on each block over time. Pedestrian(2) and business(3) data sourced from San Diego Regional Data Library. Weather information was sourced from MeteoStat(4). **Unfortunately, basic correlations proved very weak.**

Exogenous Factor	**Correlation with Rough Sleeper Counts**
Weather (temperature)	0.014
Pedestrian Traffic, Saturday Evening	0.033
Pedestrian Traffic, Wednesday Lunch	0.084
Pedestrian Traffic, Wednesday Night	0.044
Business Prevalence	-0.007

Rather than complicate our model with weak correlates, we decided to pursue learning techniques that capitalize on auto-correlative patterns in the data (the number of sleepers found previously on a block are the largest predictor for future encounters on the same block). We further elaborate on this technique in our methods section.

4 Methods

HOPE consists of an online prediction, path-optimization, and learning loop. To generate an optimal survey path at the beginning of each day, HOPE:

1. **Predicts** the number of rough sleepers at each block.

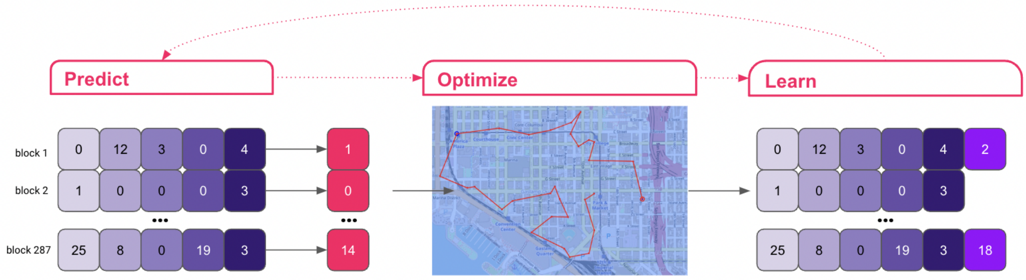

Figure 2: A visualization of the Predict-Optimize-Learn Loop.

2. **Optimizes** a block-by-block path using node rewards equivalent to expected rough sleepers at each block and constraints dictated by distance and time.
3. **Learns** by keeping track of the actual number of rough sleepers encountered at each block, and uses the updated counts for the next day's prediction step.

Step	**Goal**
Predict	$f(block\,features) \rightarrow [sleepers_{block\,1}, sleepers_{block\,2}, ...sleepers_{block\,n}]$
Optimize	$f(constraints, [sleepers_{block\,1}, ...sleepers_{block\,n}]) \rightarrow [visit\,block_1, visit\,block_2...visit\,block_v]$
Learn	$f(encountered\,sleepers) \rightarrow update(block\,features)$

Using this online system, HOPE is able to adapt to changing distributions in rough sleeping patterns and more accurately route outreach workers throughout multiple time steps.

4.1 Prediction

Several methods were explored to accurately predict the number of rough sleepers on a block-by-block basis. As discussed in the **Data Section**, there were very few co-variates that could be exploited to predict sleepers. The most salient predictor of rough sleepers at a given block was the history of sleepers found there previously. To exploit this auto-correlation in our prediction step, we explored:

1. **A heuristic** based approach, in which the last recorded sleeper count for each block is used as the next time step's prediction.
2. **Distribution Sampling**, in which we fit a distribution to the number of sleepers previously found at each block and then sample from it. Distributions were fit and sampled using the distfit python library (5).
3. **Time Series Forecasting**, in which the number of rough sleepers at each block is modeled as a time series and then a predictor fit to a training set outputs the next expected element. We experimented with two time series forecaster, N-Beats and a Temporal Fusion Transformer (TFT) as implemented by the darts python library (6). Both the N-Beats and TFT predictor were trained with a time-series input-chunk length of 7 and an output-chunk length of 1.

The input for all these predictors was the block-by-block rough sleeper counts represented in the first column of Figure 2. The output was a single number for each block representing the expected number of rough sleepers at the next time step. The predictors were all seeded using data from the first 12 out of 88 surveys on all 287 blocks. The predictors were then continuously updated as part of the online predict-optimize-learn loop during a simulation-based evaluation on the remaining 76 surveys.

4.2 Optimization

The block-by-block prediction of rough sleepers is then fed into a pathing optimization algorithm implemented in Google OR-Tools (7). Using the routing solver, we implemented the following maximization problem:

Maximize

$$max \sum_{i=0}^{blocks} sleepers_i * visited_i + !sleepers_i * 0.1 * visited_i$$

Such that

$$\sum_{v=0}^{visitedblocks} engagement(sleepers_v) + traveltime_{v-1,v} \leq 480min$$

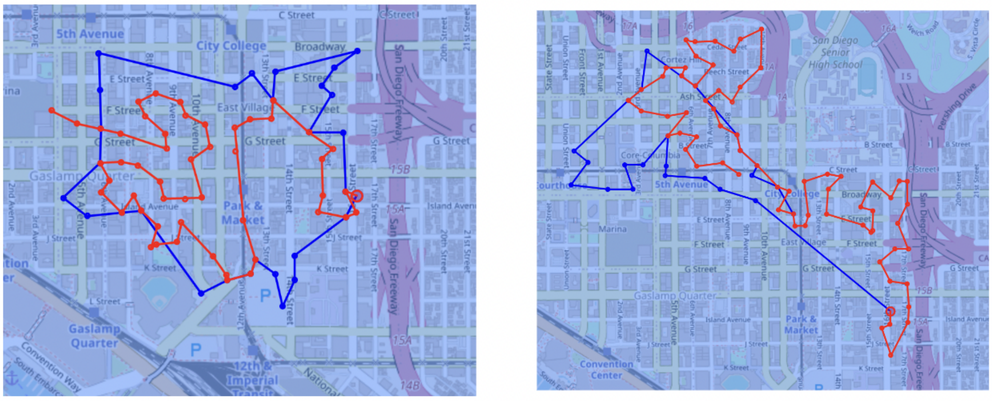

Figure 3: An example of our methodology's output on two different days using a TFT predictor, our most successful. The blue line shows the ground truth optimal path for this time period, and the red line shows the optimal path as predicted by our model. As seen, the TFT delivers good results in predicting the best path despite changing underlying distributions of rough sleepers over the two days.

where $sleepers_i$ represents predicted sleepers at $block_i$ and $visited_i = 1$ if $block_i$ is visited, 0 if not. The second term is a small reward for exploring blocks with no sleepers predicted.

where $traveltime_{v-1,v}$ is computed by dividing the geographic distance between the visited blocks by the average travel speed determined by user interviews (50m/min). Expected engagement time was based on the number of sleepers expected (5min * sleepers) with a cap of 30 minutes per block (again based on outreach worker interviews). 480 min was derived based on the average 8 hour shift length.

The starting node of the routing problem was set to the block corresponding to the San Diego Downtown Partnership's address. The path optimization algorithm outputs a path that visits each block at most once, starts and ends at the start node ($[startnode, visit_1, visit_2, ...visit_v, startnode]$) and maximizes the number of expected sleeper encounters within the 8-hour shift. A visualization of the output is shown in the middle column of Figure 2.

4.3 Learn

As outreach workers traverse the optimized path, the number of *actual* encountered sleepers is tracked at a block level. These observations are then fed back into the next predict step.

5 Evaluation

Overall, the goal of this work is to produce the optimal route for outreach workers to follow in order to encounter the greatest number of rough sleepers. Because we do not yet have a live partnership in place, **we evaluate HOPE by simulating its performance on 76 timesteps that we have ground truth data for** (recall the prediction models are seeded with the first 12 of 88 full downtown surveys). Because we have the real counts of rough sleepers for each block, we can compare the predicted route to the ground truth optimal route at each timestep. Figure 3 shows a visual of this evaluation. During the simulation, if the actual time ever exceeded 480 minutes (an 8 hour shift), that day's simulation immediately ended. Actual time was computed using the same method as in the optimization step - travel and engagement time. Not completing a full predicted path is possible because the worker may encounter more actual rough sleepers than predicted and extra time is spent engaging with those individuals, meaning not all planned nodes can be visited within 480 minutes. This ensures that all strategies are being compared over the same amount of time - 480 minutes. The ground truth optimal path is of course already optimized to this constraint.

To evaluate each of our methodologies – last-found heuristic-based, distribution sampling, and time series forecasting – we compute three primary metrics: (1) the percentage of rough sleepers encountered out of the total number possible, (2) the number of unique blocks visited, and (3) the number of days our predicted path actually encounters no rough sleepers. We also compute these metrics for two naive methods, (1) a nearest neighbor approach, where the outreach worker travels from one block to the next closest non-visited block and (2) the ground truth optimal path from the first 12 surveys of training data that encountered the maximum number of sleepers. **These naive approaches as well as the last-found heuristic are our most accurate**

	Predictor	Pathing	Total Found (%)	Avg % Found of Daily Optimal	Blocks Visited	0 Days
Optimal	Ground Truth	Optimizer	15,395 (**100%**)	100%	237	0
Naive Approaches	*no prediction*	Nearest Neighbor	9,588 (**62.3%**)	53.4%	126	16
	no prediction	Best Path From First 12	10,088 (**65.5%**)	48.8%	28	24
	Last Step Results	Optimizer	8,873 (**57.6%**)	46.07%	164	13
	Distribution sampling	Optimizer	9,147 (**59.4%**)	50%	162	5
Time Series Approaches	N-Beats	Optimizer	10,048 (**65.3%**)	54.6%	166	5
	Temporal Fusion Transformer (TFT)	Traveling Salesman	**10,817 (70.3%)**	**59.0%**	**164**	**3**

Figure 4: Chart of our results for each evaluation criteria and each prediction methodology.

representation of how outreach workers navigate today: roughly using past rewards and geographic proximity to inform their routes. By evaluating these paths as part of our simulation, we seek to compare HOPE against a real outreach worker as best we can.

6 Discussion

Given these results, we can conclude that the time series approaches, N-Beats and TFT, perform best overall. They each balance a high % of rough sleepers found, a diversity of blocks visited, and low number of days where no rough sleepers are encountered, thus indicating a more balanced and efficient approach than existing naive methods. There is value in encountering the same rough sleepers over time to establish trust, as well as encountering new rough sleepers to introduce them to outreach services. We believe that a balanced approach helps pursue both goals. Given these goals, the TFT predictor coupled with the route optimizer performed exceptionally well and gives us reason to believe that such a method could be a valuable addition to outreach workers' toolkits.

One known limitation is that, due to privacy concerns, this data does not track any specific individuals. As a result, we can make no conclusions about the true ratio of new to previously encountered individuals that outreach workers engage with through Project HOPE's pathing. However, this is also a valid trade-off accepted to preserve the privacy of all individuals. Future data collection methods could improve upon this by introducing a binary flag or count that outreach workers can use to indicate if any given rough sleepers they encounter on a block are new to the program or not. Such an approach would preserve privacy while allowing for outreach efforts to find its ideal balance between new and already-encountered individuals.

Additionally, we recognize potential ethical concerns of rough sleeper location predictions being misused in ways that could harm the very individuals we are intending to help (policing, harassment, etc.). For this reason, future work to fully implement a system like Project HOPE must be safeguarded to ensure it is only being used as intended by outreach workers.

7 Future Work

Several items of future work remain for the vision of Project HOPE to be fully realized. First, although we were initially informed of this problem through conversations with an outreach worker based in New York City, we were unable to communicate with anyone working in San Diego. Gathering specific context from our end users remains a large gap that could influence its usability. Furthermore, our system currently does not adapt to mid-day changes. For example, if an outreach worker travels to an unexpected block or an engagement takes far more time than expected, HOPE will need to be rerun. HOPE runs in a matter of seconds, so adding this functionality is not difficult - but is essential before a minimum viable product can be deployed into real-world usage. Finally, while we were not able to confirm any strong exogenous correlates with the San Diego data, there is reason to believe that factors such as weather, crime, infrastructure, and traffic could significantly enhance HOPE's predictive power in other locations. Unfortunately, the type of data collection used in San Diego has not be copied in other locations.

Given the promise of HOPE's prediction mechanism, we believe that the introduction of an application targeted at outreach workers to allow them to digitally collect data while on their shifts would be incredibly beneficial for this problem. If such a data collection infrastructure existed, we could begin experimenting with other potential prediction mechanisms to further improve the efficacy of HOPE in additional locations. A likely requirement for launching would be for us to construct a simple mobile UX to enable this data collection ourselves as none currently exist.

8 Contributions and Acknowledgments

We are thankful for the CS 288: AI for Social Impact class at Harvard University. The course introduced us to the field of AI for Social Impact and its related discussions. In particular, we were encouraged to see a healthy discussion of stakeholder engagement and ethical tradeoffs. We would like to acknowledge and thank Professor Milind Tambe as well as the Teaching Assistants, Sonja Johnson-Yu and Paula Rodriguez Diaz.

As for our contributions, Noah undertook the heavy lifting and performed the bulk of the AI prediction and routing steps as well as the outreach worker interviews. Elena led the context, discussion, and evaluation, and Neeraj contributed the naive approaches.

References

Unhoused Care – Downtown San Diego Partnership. https://downtownsandiego.org/clean-and-safe/unhoused-care/.

San Diego Regional Data Library. "San Diego Regional Library Data Repository." Pedestrian Density By Block and Time, 13 June 2019, https://data.sandiegodata.org/dataset/sandiegodata-org-pedestrian-blocks/.

San Diego Regional Data Library. "San Diego Regional Library Data Repository." San Diego County Business Sites, 13 June 2019, data.sandiegodata.org/dataset/sangis-org-business-sites. Accessed 2 Dec. 2022.

"MeteoStat." PyPi, 17 Aug 2022, https://pypi.org/project/meteostat/

"Distfit." PyPI, 10 May 2022, https://pypi.org/project/distfit/.

Forecasting Models — Darts Documentation, https://unit8co.github.io/darts/generated_api/darts.models.forecasting.html.

"Vehicle Routing, OR Tools; Google Developers." Google, Google, 4 Oct. 2022, https://developers.google.com/optimization/routing.

Busboom, Eric. "Downtown-Partnership/Homelessness Space and Time.ipynb at Master · Sandiegodata-Projects/Downtown-Partnership." GitHub, Downtown San Diego Partnership, 24 Oct. 2019, https://data.sandiegodata.org/dataset/sandiegodata-org-dowtown-homeless/.

VanBerlo, Blake, et al. "Interpretable machine learning approaches to prediction of chronic homelessness." Engineering Applications of Artificial Intelligence 102 (2021): 104243.

Shah, Ovais Rehman, Lori Willoughby, and Nicholas Bowersox. "Tackling homelessness through AI powered social innovations: A novel and ground-breaking assessment of criminal victimization of homeless populations in los angeles employing predictive analytics and machine learning models such as ARIMA and LSTM." Issues in Information Systems 22.3 (2021).

Wilde, Harrison, et al. "A Recommendation and Risk Classification System for Connecting Rough Sleepers to Essential Outreach Services." Data & Policy, 3.2 (2021).

Guevara, Cesar and Santos Peñas, Matilde. (2021). Crime Prediction for Patrol Routes Generation Using Machine Learning. 10.1007/978-3-030-57805-3_10.