

C2 SMART

CONNECTED CITIES WITH
SMART TRANSPORTATION



A USDOT University Transportation Center

New York University

Rutgers University

University of Washington

University of Texas at El Paso

The City College of New York

City-scalable Destination Recommender System for On- demand Senior Mobility

August 2018



City-scalable Destination Recommender System for On-Demand Senior Mobility

Joseph Y. J. Chow
New York University

Assel Dmitriyeva
New York University

Daniel Fay
New York University

C2SMART Center is a USDOT Tier 1 University Transportation Center taking on some of today's most pressing urban mobility challenges. Using cities as living laboratories, the center examines transportation problems and field tests novel solutions that draw on unprecedented recent advances in communication and smart technologies. Its research activities are focused on three key areas: Urban Mobility and Connected Citizens; Urban Analytics for Smart Cities; and Resilient, Secure, and Smart Transportation Infrastructure.

Some of the key areas C2SMART is focusing on include:

Disruptive Technologies

We are developing innovative solutions that focus on emerging disruptive technologies and their impacts on transportation systems. Our aim is to accelerate technology transfer from the research phase to the real world.

Unconventional Big Data Applications

C2SMART is working to make it possible to safely share data from field tests and non-traditional sensing technologies so that decision-makers can address a wide range of urban mobility problems with the best information available to them.

Impactful Engagement

The center aims to overcome institutional barriers to innovation and hear and meet the needs of city and state stakeholders, including government agencies, policy makers, the private sector, non-profit organizations, and entrepreneurs.

Forward-thinking Training and Development

As an academic institution, we are dedicated to training the workforce of tomorrow to deal with new mobility problems in ways that are not covered in existing transportation curricula.

Led by the New York University Tandon School of Engineering, C2SMART is a consortium of five leading research universities, including Rutgers University, University of Washington, University of Texas at El Paso, and The City College of New York.

c2smart.engineering.nyu.edu

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

Acknowledgements

Xuebo Lai helped with early development of the recommender system which is gratefully acknowledged. We would also like to thank David Sinclair, a student at NYU Wagner, for his help in conducting the elderly mobility survey in NYC. The survey research was conducted in collaboration with Professor Kelvin Cheu at the University of Texas at El Paso and his team there. Participation in the NSF-RCN Workshop on Smart and Connected Communities and Aging Population at Stony Brook University on April 20, 2018 helped shed light on some recent senior surveys conducted nationwide, which is much appreciated.

Executive Summary

Mobility-on-demand (MOD) services—rideshare, car- and bike-share, e-hail taxis, microtransit, and more—have been on the rise due to advances in information and communications technologies (ICTs). Despite their increasing popularity, many operational challenges make them hard to sustain. Operators need to consider more sophisticated dynamic operations. One of the biggest contributing factors to high operating costs is the presence of incidents that disrupt the planned route: non-recurring traffic delays; customer schedule delays, reschedules, or cancellations; or cancellation, closure, or malfunction at a destination such as malfunction of an access elevator at a transit station. MOD services can be smarter by interacting with users and recommending activity destinations to them. This is also a first step in making mobility companies act as physical search engines for travelers.

What is needed is an efficient learning mechanism for MOD services so that destinations they recommend can help the service efficiently learn the users' preferences over time. The learning problem of selecting options under repeated trials of those options is called a contextual bandit problem. The fundamental trade-off is to balance an option selection that could efficiently learn the uncertainty (exploration) associated with that option while still providing users with rewarding choices (exploitation). A good learning mechanism is one that, over multiple trials, minimizes its expected regret. Destination recommender systems conducted by a MOD service are different from conventional recommender systems. There is a fundamental conflict between trying to minimize regret by learning as efficiently as possible and trying to minimize operating costs. A destination may be highly rated and offer a good learning opportunity for a user but having to re-route the MOD service vehicle (and its passengers) to serve the location might significantly increase operating costs.

We set out to accomplish three objectives in this project:

1. Better understand the mobility needs of the elderly across different cities. This is accomplished via a joint survey conducted with collaborators from the University of Texas, El Paso (UTEP), on elderly living in El Paso, TX, and in New York, NY.
2. Implement a proof-of-concept of a recommender system that can be readily adapted to MOD services, one that considers routing constraints.
3. Conduct computational experiments with the proof of concept to demonstrate the existence of the effect that adding spatial constraints has on the performance of a recommender system. Based on these computational experiments, we draw new guidelines for expanding this research for MOD service providers using publicly available data.

The research on destination recommendation systems has paid little attention to the integration of contextual information with the recommendation algorithm. The most related problem setting to ours is Brunato's PILGRIM: A Location Broker and Mobility-Aware Recommendation System.

Two students from NYU, with proper IRB certification to survey human subjects, conducted a survey at senior centers. The questionnaire (see final report from UTEP Year 1 project) was conducted at five senior centers in New York City (NYC) from February 5 to February 16, 2018. A total of 61 responses were received.

With the survey results from El Paso and NYC, it is possible to compare elderly mobility responses across the two different cities. This helps provide insights on how these preferences scale from El Paso (population 680K in 2016) to NYC (8.5 million in 2016). Our survey provided a better understanding of these needs and preferences: in cities, the elderly tend to more frequently use smartphone devices, and uncertainties associated with a trip are the top concerns when traveling (cost, weather, on-time reliability). Accessibility is also a major challenge that needs to be overcome. These findings provide motivation for improving the intelligence of MOD services to better cater to the elderly among other population groups.

The second objective was to implement a prototype recommender system that can be adopted by MOD services. A code was developed based on the UCBGLM algorithm from Li et al. (2010,2017) with a modification to the payoff function to include increase in routing cost as an additional variable. The code is available on our public repository.

The third objective was to conduct some computational tests with the prototype system to evaluate our hypothesis that it is important to explicitly incorporate routing constraints and that these constraints will tend to worsen the performance of the algorithm. Despite running into data security issues during the project (as reported in our final quarterly progress report), we were able to conduct some preliminary assessments that suggest the recommender system's value increases when there is a greater need for it.

Table of Contents

Executive Summary.....	iv
Table of Contents.....	vi
List of Figures.....	vii
List of Tables.....	viii
1. Introduction.....	1
2. Literature Review.....	4
2.1. Elder mobility.....	4
2.2. Recommender systems based on contextual bandit problems.....	5
3. Elderly Mobility Survey.....	8
3.1. Survey implementation.....	8
3.2. Survey results.....	9
3.3. City-scalability comparison between El Paso and NYC elderly mobility.....	11
4. Routing-Constrained Recommender System Proof-of-Concept.....	15
4.1. Recommendation engine.....	15
4.2. Routing subproblem.....	18
5. Computational experiments.....	23
5.1. Experimental setup.....	23
5.2. Results.....	25
6. City-wide deployment, deliverables, and technology transfer.....	29
6.1. Deployment.....	29
6.2. Deliverables.....	30
6.3. Technology transfer.....	31
7. Conclusion.....	32
References.....	34

List of Figures

Figure 1: Consideration of destination recommendation in a routing context. 2

Figure 2: Categories of location-based recommendation systems (Rehman et al., 2017). 6

Figure 3: Senior centers surveyed. 8

Figure 4: Electronic device usage. (source: Cheu et al., 2018) 12

Figure 5: Concerns when planning a trip. (source: Cheu et al., 2018) 12

Figure 6: Mobility challenges. (source: Cheu et al., 2018)..... 13

Figure 7: Desired app functions. (source: Cheu et al., 2018)..... 14

Figure 8: Recommendation engine framework. 15

Figure 9: Heatmap of establishments with Yelp Open Data in Las Vegas..... 17

Figure 10: Average regret trajectory without routing cost consideration..... 18

Figure 11: Illustration of complexity of routing subproblem. 19

Figure 12: Insertion heuristic solution..... 22

Figure 13: Simulated ratings (out of 5) across the 100 zones. 25

Figure 14: Realized recommendations for base scenario. 26

Figure 15: Distribution of recommended zones after (left) 500 trials, (right) 2000 trials..... 26

Figure 16: Average regret comparison. 27

Figure 17: Spatial distribution of recommendations when $tt = 4/3$ base..... 28

List of Tables

Table 1: New York Survey Trip Frequencies..... 10

Table 2: New York Survey Modes of Transportation 10

Table 3: Raw Features from Yelp..... 30

1. Introduction

Mobility-on-demand (MOD) services—rideshare, car- and bike-share, e-hail taxis, microtransit, and more—have been on the rise (Deloitte, 2015) in the last few years due to advances in information and communications technologies (ICTs). Despite their increasing popularity, many operational challenges make them hard to sustain. For example, paratransit service in New York costs the public transit agency as much as \$71 per trip to operate (Crains, 2016). Due to high costs of operation, various services have recently failed as well: Car2Go in San Diego, Kutsuplus microtransit in Helsinki, Bridj microtransit, bike-share in Seattle, among others.

To prevent the potential demise of MOD services, operators need to consider more sophisticated dynamic operations. One of the biggest contributing factors to high operating costs is the presence of incidents that disrupt the planned route: non-recurring traffic delays; customer schedule delays, reschedules, or cancellations; or cancellation, closure, or malfunction at a destination such as malfunction of an access elevator at a transit station. The impacts of these events have led to special provisions by paratransit agencies on conditions for canceling or rescheduling appointments, for on-demand scheduling software providers like GIRO to incorporate specific features to respond to cancellations, added trips, or en-route delays, and for researchers to propose recourse strategies (e.g. Vodopivec et al., 2015).

MOD services can be smarter by interacting with users and recommending activity destinations (e.g. Chow and Liu, 2012) to them. Like services offered by companies like Amazon and Netflix, an MOD service can suggest destinations to a user when they are booking a trip or when an incident occurs that would either significantly increase the cost of delivering a passenger or do so at a higher risk of delaying other passengers to their appointments. This would also allow services like Uber, Lyft, Didi, and Via to offer an option for a traveler to book a trip to “a restaurant”, for example, and leave it to the service to recommend specific nearby restaurants to the user that they are likely to enjoy. This is also a first step in making mobility companies act as *physical search engines* for travelers.

Recommender systems are not new (see Adomavicius and Tuzhilin, 2005). What makes them most effective is when there is abundant data about personal preferences for different options from which to make forecasts. For example, given enough data, one can use singular value decomposition to break down observations of users’ preferences into common categories and preferences within them (Li et al., 2010). The problem, particularly with new services with limited initial data, is how to build up the database in an efficient manner. While some services may simply survey their users (e.g. Netflix asking users to rate movies), it would not be efficient for MOD services to do the same because destination recommendations are highly contextual and depend on the location of the user and a vehicle’s route.

What is needed is an efficient learning mechanism for MOD services so that destinations they recommend can help the service efficiently learn the users' preferences over time.

The learning problem of selecting options under repeated trials of those options is called a contextual bandit problem (Li et al., 2010). The fundamental trade-off is to balance an option selection that could efficiently learn the uncertainty (exploration) associated with that option while still providing users with rewarding choices (exploitation). A good learning mechanism is one that, over multiple trials, minimizes its expected regret.

Learning considerations in travel behavior research tend to focus on modeling travelers' learning (e.g. Arentze and Timmermans, 2003). There are also efforts to apply learning algorithms in mobility services like taxis (e.g. Phithakkitnukoon et al., 2010), but no work has yet been applied to develop and understand user destination recommender systems for MOD services.

Destination recommender systems conducted by a MOD service are different from conventional recommender systems. There is a fundamental conflict between trying to minimize regret by learning as efficiently as possible and trying to minimize operating costs. A destination may be highly rated and offers a good learning opportunity for a user but having to re-route the MOD service vehicle (and its passengers) to serve the location might significantly increase operating costs. This is illustrated in Figure 1. In the left-hand side, a MOD shuttle currently serves two green passenger waypoints when a new customer enters the system with pickup and drop-off locations. This is the state of practice. In the "smart mobility" scheme, the MOD shuttle may instead recommend the new passenger to either the red or yellow destinations as shown on the right-hand side. The red location may end up having a lower review/rating score from the public, but recommending the passenger to that destination would cost the shuttle less in routing cost (as indicated by the length of the red arrows compared to the yellow arrows).

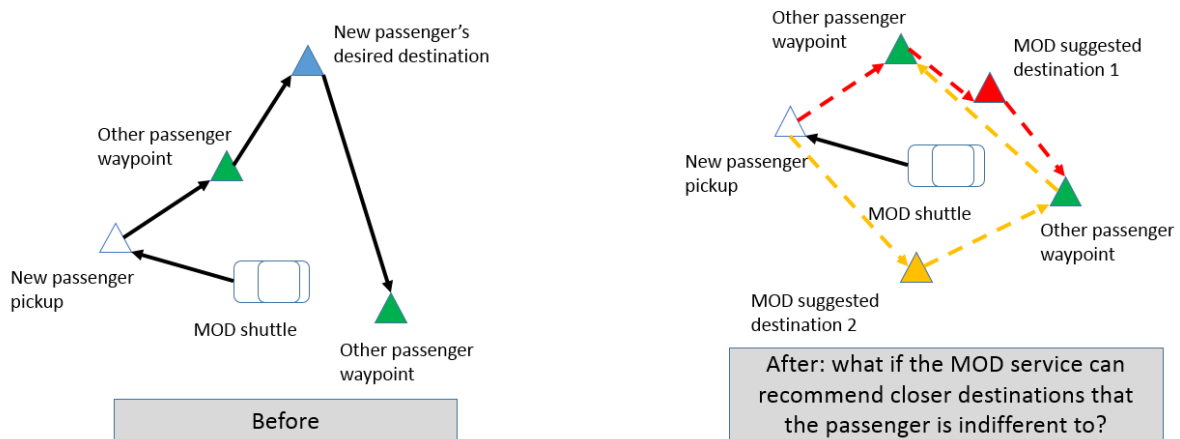


Figure 1: Consideration of destination recommendation in a routing context.

We hypothesize that there is a non-trivial spatial effect on the regret minimization bound, and due to differences in spatial distribution of destinations and road infrastructure in different cities, this effect will vary from city to city.

We set out to accomplish three objectives in this project:

1. Better understand the mobility needs of the elderly across different cities. This is accomplished via a joint survey conducted with collaborators from Professor Kelvin Cheu's team at the University of Texas, El Paso (UTEP), on elderly living in El Paso, TX, and in New York, NY.
2. Implement a proof-of-concept of a recommender system that can be readily adapted to MOD services, one that considers routing constraints.
3. Conduct computational experiments with the proof of concept to demonstrate the existence of the effect that adding spatial constraints has on the performance of a recommender system. Based on these computational experiments, we draw new guidelines for expanding this research for MOD service providers using publicly available data like Yelp (2017).

2. Literature Review

The literature review is divided into two subsections, one describing the literature on the social need for these systems and motivating the survey conducted; the second describing the literature on other studies related to recommender systems leading to the system design used for our prototype. The elderly mobility literature review is also included in the joint paper Cheu et al. (2018).

2.1. Elder mobility

The U.S. Census Bureau defines a senior as a person over the age of 65 (U.S. Census, 2017), with approximately 50 million seniors in the U.S. in 2015. This translates to approximately 12% of all trips and 10% of all miles traveled in the U.S. in 2009 taken by elders (AARP, 2015). Accessibility to transportation options is a major challenge for seniors as transportation systems are inadequately designed to meet their needs. Alsnih and Hensher (2003) highlight these issues and suggest public transport solutions. To be more specific, a number of problems associated with aging contribute to increased mobility problems, including arthritis (March et al., 1998), dementia (Geerlings et al., 1999), and disorientation (Hanley, 1981).

Surveys have been conducted on senior mobility. For example, AARP Inc. commissioned the Understanding Senior Transportation Survey in 1998 (Ritter, 2002). The Independent Transportation Network (ITN) (Freund and McKnight, 1997), founded as part of the Transit IDEA program, has sample data ($n = 2094$) across the U.S. relating senior demographics to travel mode preferences (Bird et al., 2017). Various cities and counties have conducted their own surveys (e.g. Sarasota County, CA (SCOPE, 2016)). Silvis and Niemeier (2009) noted that ridesharing tends to be the second most common transport mode for seniors behind driving. They conducted a survey in California retirement homes and concluded that seniors with more active social networks use rideshare more regularly.

Despite extensive prior surveys on elderly mobility, the insights on their preferences for MOD services and related topic areas (e.g. smartphone use, data privacy, shared mobility adoption) have been limited.

2.2. Recommender systems based on contextual bandit problems

The multi-armed bandit problem is a sequential optimization problem in which one of a set of options (the “arms”) is chosen each trial in which the reward distribution starts out unknown. Each selection of an arm allows the system to learn more about the distribution of the reward for that arm but also takes on the reward of that arm in that trial. As a result, there is a fundamental trade-off, over a sequence of trials, to select arms that provide the best rewards (exploitation) against the need to sample different arms (exploration). It is a type of dynamic decision process in which the system would ideally start off with more exploration so that the sacrifices in rewards early on will benefit the choices in later trials in a sequence. Over a sequence of trials, the measure of performance is typically quantified in terms of “regret”. Generally speaking, regret is a measure of the difference between the reward a system obtains upon deciding versus the best reward that the system can get. There is experienced regret if the true rewards are known, but since in practice this is typically unknown, the anticipated or expected regret is computed.

Different learning algorithms can be used to decide which arms to select such that, in worst case scenarios, the accumulated expected regret should not exceed a certain bound.

A variant of the multi-armed bandit problem is the contextual bandit problem, where the system explicitly obtains feedback from the user (they will accept or reject a recommendation), and that information is then used to update a model of contextual information of the users and the attributes of their options. In other words, the user’s preference for options are explicitly modeled. The objective in such a system is to maximize the number of acceptances of the recommendations over a sequence of trials.

The contextual bandit learning algorithm has a basic formulation which we provide an overview from Li et al. (2010). There is a sequence of trials $t = 1, 2, \dots$, wherein each trial t :

1. An algorithm observes a user u_t and a set \mathcal{A}_t of options available in that trial. Each option has a set of attributes associated with the options pertaining to the user, $x_{ta}, a \in \mathcal{A}_t$. This is the “context”.
2. Based on observed rewards in previous trials, the algorithm chooses an arm $a_t \in \mathcal{A}_t$ and receives a random reward r_{ta_t} whose expectation depends on user u_t and chosen arm a_t .
3. The algorithm then improves the arm selection strategy with a new observation $(x_{ta_t}, a_t, r_{ta_t})$, with no observation of unchosen arms.

The total T -trial payoff is defined as $\sum_{t=1}^T r_{ta_t}$. The optimal expected T -trial payoff is $E[\sum_{t=1}^T r_{ta_t^*}]$, where a_t^* is the arm with maximum expected payoff at trial t . The T -trial regret $R_A(T)$ for an algorithm A is defined in Eq. (1).

$$R_A(T) \equiv E \left[\sum_{t=1}^T r_{ta_t^*} \right] - E \left[\sum_{t=1}^T r_{ta_t} \right] \quad (1)$$

With rapid growth of web applications, contextual bandits have many applications in recommendation and web-based advertising (Agarwal et al., 2009; Li et al., 2010). The most studied model in the contextual bandit literature is the linear model (e.g. Auer, 2002; Rusmevichientong and Tsitsiklis, 2010; Yasin Abbasi-Yadkori et al., 2011; Chu et al., 2011). Several studies have also looked at stochastic generalized linear models of which logistic regression is a special case for contextual bandit problems (Filippi et al., 2010; Jun et al., 2017; Li et al., 2017).

Destination recommendation can be viewed within a contextual bandit learning setting. For example, there are many locations in which there is no information known about how much a user prefers one location over another. Such information is scarce and typically unobservable. A user's preference for a recommended location can depend on a set of attributes x_{ta} like establishment rating, type of establishment, proximity to the user's work or home, among others. The payoff can be whether or not the user accepts the recommendation.

The research on destination recommendation systems has paid little attention to the integration of contextual information with the recommendation algorithm. Rehman et al. (2017) conducted an extensive survey on location-based recommendation systems (LBRS). According to their classification of LBRS in Figure 2, no effort has been made towards incorporating routing constraints in the destination recommendations.

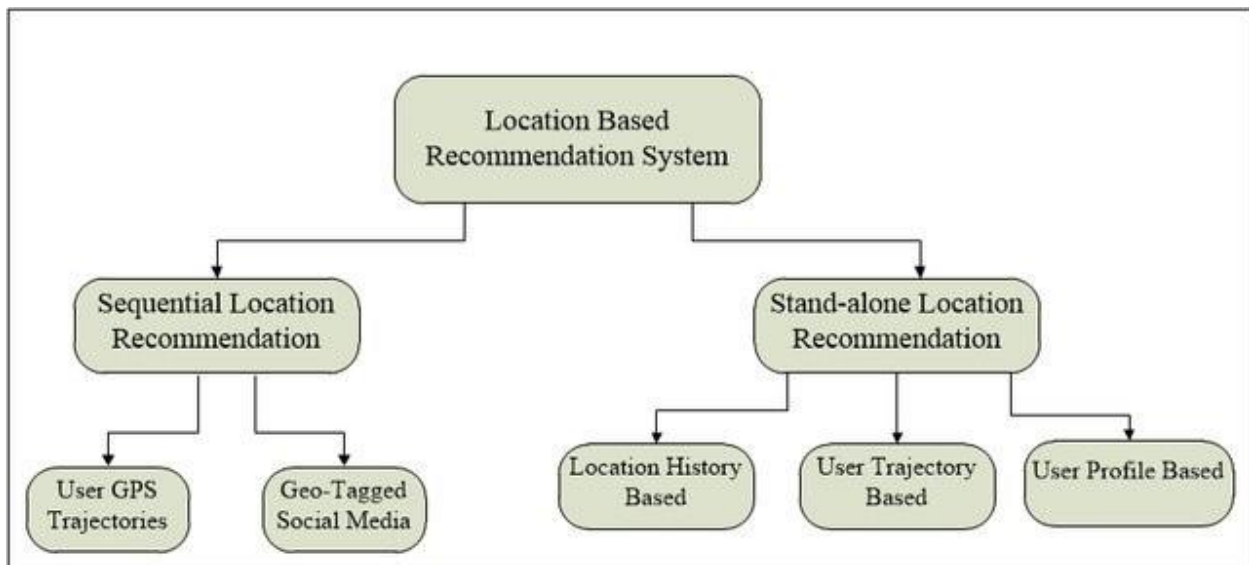


Figure 2: Categories of location-based recommendation systems (Rehman et al., 2017).

By connecting the physical world to a LBRS, unique properties of locations bring new challenges such as 1) location context awareness, 2) the heterogeneous domain, and 3) the rate of growth (Bao et al., 2015).

The most related problem setting to ours is Brunato's PILGRIM: A Location Broker and Mobility-Aware Recommendation System (Brunato and Battiti, 2003). They used the user's position as a relevant piece of information while selecting and ranking links of interest to the user. The authors propose a mobility-aware recommendation system that uses the location of the user to filter recommended links. Their work designed a middleware layer where the location broker maintains a historic database of location and corresponding links used in the past.

3. Elderly Mobility Survey

This section describes the data collection efforts in New York conducted by this team in a joint effort with UTEP to gain a better insight on elderly mobility needs in the context of MOD and smart cities-oriented setting.

3.1. Survey implementation

Two students from NYU, with proper IRB certification to survey human subjects, conducted a survey at senior centers. The questionnaire (see final report from UTEP Year 1 project) was conducted at several senior centers in New York City (NYC). Based on NYCHA Facilities and Service Centers data from December 2012 (NYCHA, 2012), there are 116 occupied senior centers in NYC. Facilities are sponsored by different agencies and geographically located in the Bronx, Brooklyn, Manhattan, Queens and Staten Island boroughs. Through contacts made available with Manhattan senior centers, it was decided to focus the surveys on these locations.

After IRB approval was obtained on January 4, 2018, coordination efforts with each senior center began regarding permission to conduct the survey. Further discussions took place with the staff at each center regarding logistics (i.e. best days and times to conduct the survey). The surveys took place at five different senior centers (Figure 3) from February 5 to February 16, 2018.



Figure 3: Senior centers surveyed.

Although the survey was created in Qualtrics, the same questions were copied in Microsoft Word. Then, hard copies were brought to the senior centers to have the seniors fill out. Each completed survey form was manually uploaded to Qualtrics at the end of each survey day.

3.2. Survey results

In contrast to the survey conducted in El Paso, the survey in New York City was conducted only in English. A total of 61 responses were received. The first part of the survey recorded the standard demographic profiles of the participants. The results are highlighted below.

- 76% of participants were over the age of 65.
- 79% of participants were female.
- 93% of participants were retired.
- 60% of participants were African American, followed by 35% Hispanic or Latino.
- 75% speak English with confidence, followed by 43% who speak Spanish.
- Zip code breakout shows that East Harlem, Harlem, Hamilton Heights, and Washington Heights neighborhoods were represented.

A majority of the participants (97%) reported that they lived in an apartment, which makes sense for NYC (NMHC, 2015). Almost half (46%) of the participants reported that they did not have any impairments and/or disabilities. Of the remaining 54% who reported having impairments and/or disabilities, the three frequently reported issues were difficulty with walking, followed by visual and muscle control impairments. Approximately one third of participants (34%) reported that they did not require any assistance, followed by those who require only a cane (31%).

The frequency of destinations traveled (per week) are tabulated in Table 1. The most visited places are senior center, library, park, or gym. The frequency of use of transportation modes are tabulated in Table 2. We observe that 80% never use paratransit service, and only 13% use ridesharing service.

Table 1: New York Survey Trip Frequencies

Destination	Frequency
Volunteering Place	Never (69%)
Family Member, Relative, or Friend	Never (41%)
Grocery, Market, or Retail Shop	1 to 3 times per week (38%)
Healthcare Facility, or Pharmacy	Never (33%)
Senior Center, Library, Park, or Gym	3 to 6 times per week (44%)
Civic or Religious Center	Never (43%)
Restaurant, Coffee Shop, Diner	Never (48%)
Bank, ATM, or offices	Less than once per week (36%)

Table 2: New York Survey Modes of Transportation

Mode of Transportation	Frequency
Walking More Than ¼ Mile	7 times per week or more (30%)
Bicycle	Never (93%)
Motorcycle/Scooter	Never (100%)
Car (as Driver)	Never (90%)
Car (as Passenger)	Never (54%)
Carpool (as Driver or Passenger)	Never (87%)
Public Bus	Less than once per week (26%)
Special Bus (e.g. Lift)	Never (80%)
Taxi	Never (56%)
Rideshare (e.g. Uber, Lyft)	Never (87%)

One of the last questions asked the participants to list the factor that would motivate them the most to use a smartphone application specifically designed for their mobility needs. Among those who owned a smartphone, the most popular answers were getting to a destination efficiently, followed by the simplicity of the application. Over half (56%) of participants reported that they would not be willing to anonymously share their data collected via the application.

3.3. City-scalability comparison between El Paso and NYC elderly mobility

With the survey results from El Paso and NYC, it is possible to compare elderly mobility responses across the two different cities. This helps provide insights on how these preferences scale from El Paso (population 680K in 2016) to NYC (8.5 million in 2016).

For both surveys conducted, the smartphone was reported as the most frequently used electronic device at 49% and 62% for El Paso and New York City, respectively. Both are higher than the U.S. national average which is in the 30-40% range (Berenguer et al., 2017). This is likely due to the correlation with city residents in the two survey areas compared to rural residents in most of the country. The basic home phone was the second most-frequently used electronic device for both surveys conducted at 39% and 56% for El Paso and New York City, respectively. The results for all electronic devices used from both surveys are presented in Figure 4. The follow-up question asked the participants if they require any assistance to use such devices. Almost three-quarters (74% and 70% for El Paso and New York City, respectively) of respondents reported that they did not require assistance; however, an informal discussion with the respondents provided insights that they use their smartphone exclusively to call relatives.

In the El Paso survey, the most frequently selected concerns while making a trip within the city were on-time departure, followed by protection from extreme weather, and cost. In the New York City survey, the most frequently selected concerns while making a trip within the city were cost, followed by protection from extreme weather, and on-time departure. The top three concerns from both surveys were the same, just in different orders. The results for all types of concern when planning a trip is presented in Figure 5.

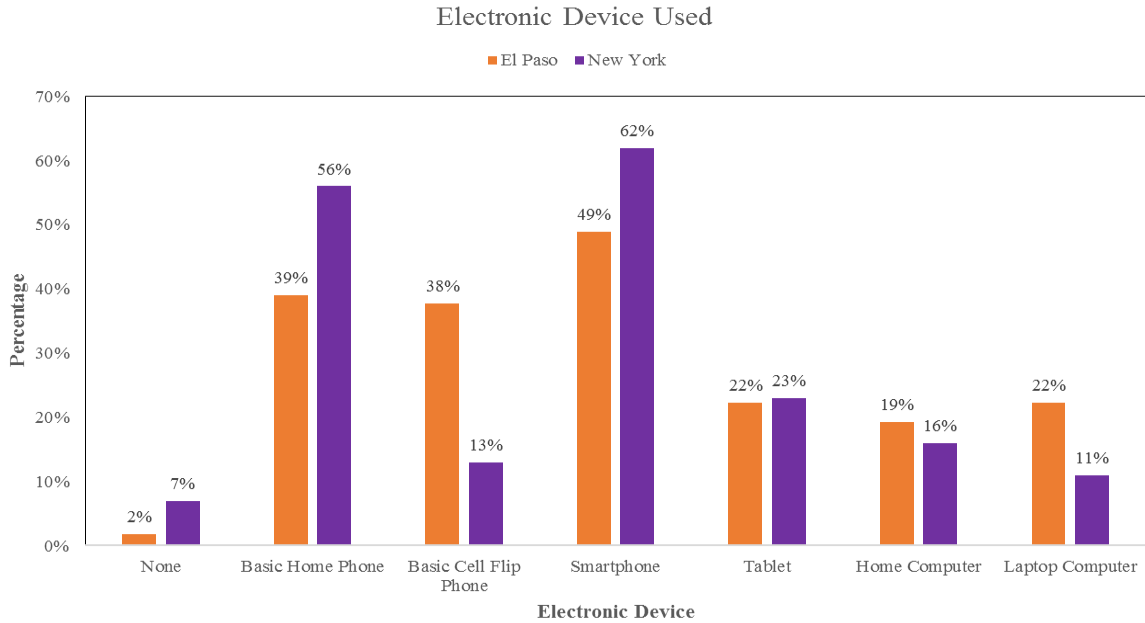


Figure 4: Electronic device usage. (source: *Cheu et al., 2018*)

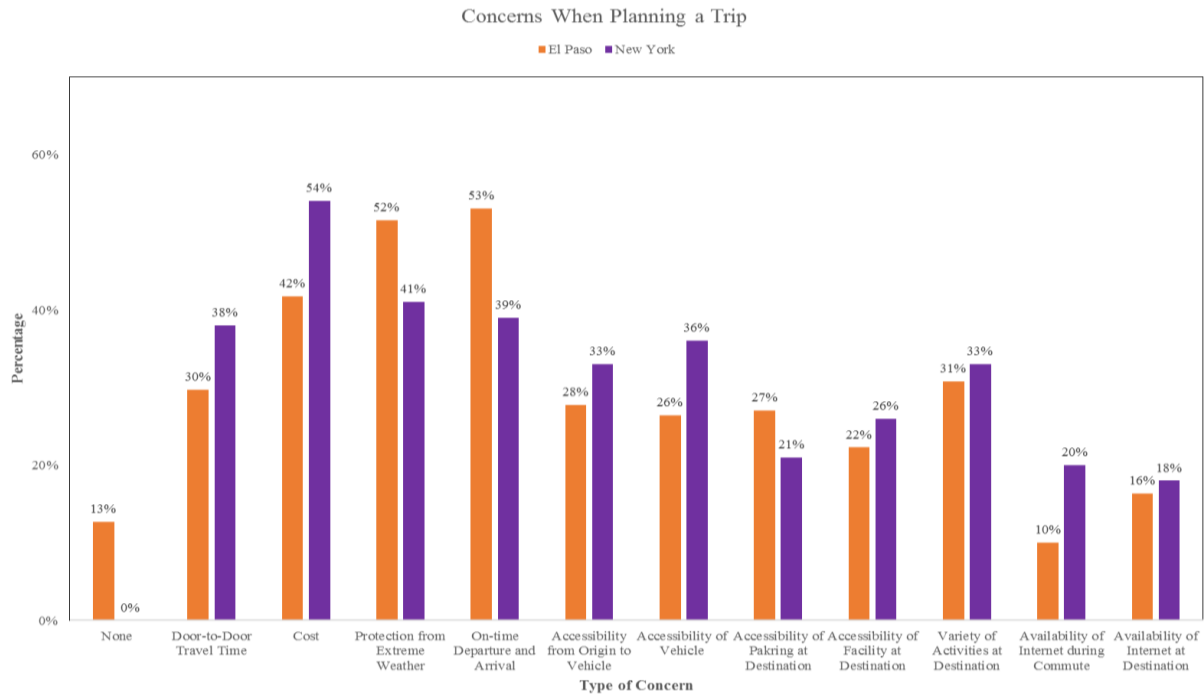


Figure 5: Concerns when planning a trip. (source: *Cheu et al., 2018*)

An open-ended question asked the participants to describe their biggest challenge when they commute in the city. In the El Paso survey, a majority of participants (62%) reported that traffic, parking and

construction was the greatest challenge when commuting in the city, followed by difficulty in walking (15%). In the New York City survey, approximately one-third of participants (33%) reported that waiting time was the greatest challenge when commuting in the city, followed by traffic, parking, and construction (19%) and accessibility to vehicles and buildings (19%). This provides justification that seniors are in direct need of a mobile application that guides them before and during their trips within the city (e.g., avoiding traffic congestion, help finding parking, avoiding construction, etc.). The results for all recorded mobility challenges are presented in Figure 6.

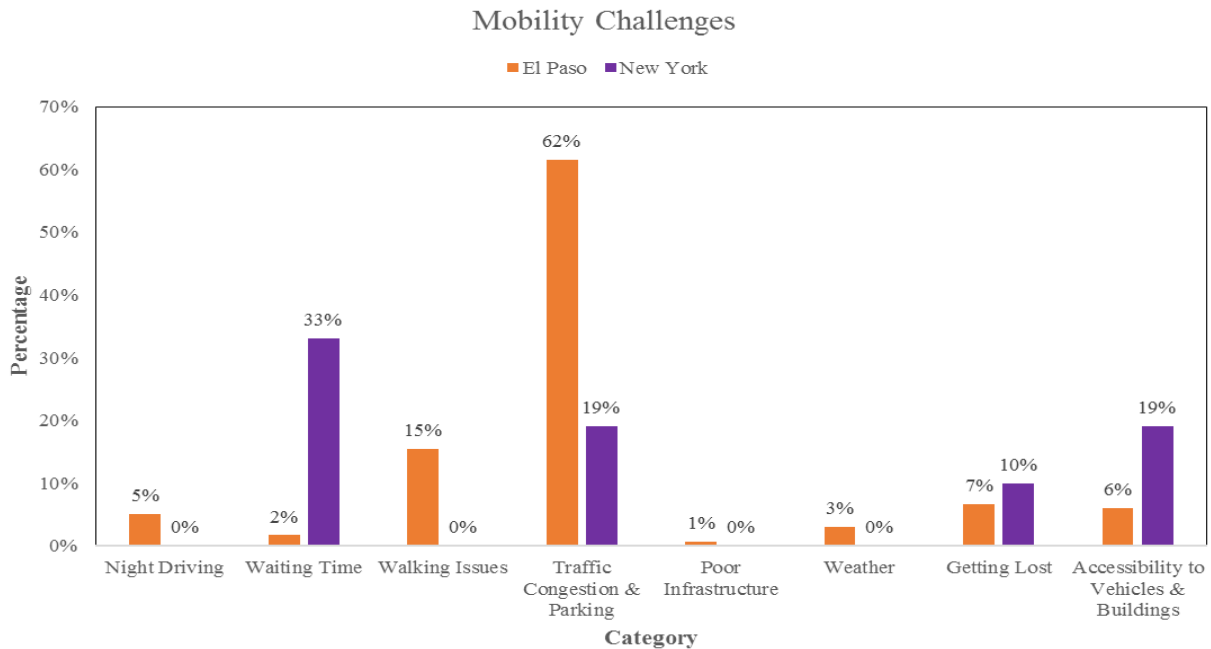


Figure 6: Mobility challenges. (source: Cheu et al., 2018)

Another open-ended question asked the participants to list one function they would like to see in a smartphone application. In the El Paso survey, the most popular answer was navigation (53%), followed by the inclusion of bus routes (20%), and then the overall simplicity and intuitiveness of the application (13%). In the New York City survey, one answer choice that was not seen in the El Paso survey was the most popular answer: adding more features for the Apps (60%). The next most popular answers were navigation (13%) and bus routes (13%). The results for all recorded desired functions are presented in Figure 7.

The survey results demonstrate how the elderly have different preferences from NYC and El Paso. One common insight was that public transit information and navigation are key features desired in smartphone apps due to the desire to get to destinations efficiently.

Desired App Functions




Figure 7: Desired app functions. (source: Cheu et al., 2018)

4. Routing-Constrained Recommender System Proof-of-Concept

Similar to Brunato and Battiti’s (2003) PILGRIM system, we separate the system into two main parts: 1) Recommendation and 2) Route Generation parts. The problem of recommending destinations is formulated as a contextual multi-armed bandit. The route generation part is modeled as a dynamic traveling salesman problem with pickups and drop-offs (TSPPD). Recommendation and Route generation parts can be loosely coupled: the results of the recommendation are the input (destinations) of the route generation. Such a loose coupling lacks the flexibility needed in many special situations, especially when errors occur, or the user behaves in an unpredicted manner. To address these problems, closer coupling of recommendation and route generation must be implemented. Close coupling results in a theoretical problem formulation.

4.1. Recommendation engine

The general framework of the proposed system is shown in Figure 8. There is a procedure to obtain a payoff matrix for users, a procedure to select a destination to recommend, and a procedure to update the database of payoffs. Based on an online database of users’ ratings for different destinations, the system uses Singular Value Decomposition (SVD) to predict ratings of destinations for each user. This results in a set of candidate destinations, \mathcal{A}_t and a d -dimensional SVD-based feature vector x_{ta} for each destination (routing cost increase is treated separately). From these values the system maintains an inverse covariance matrix A_a and a vector b_a of ratings corresponding to those features.

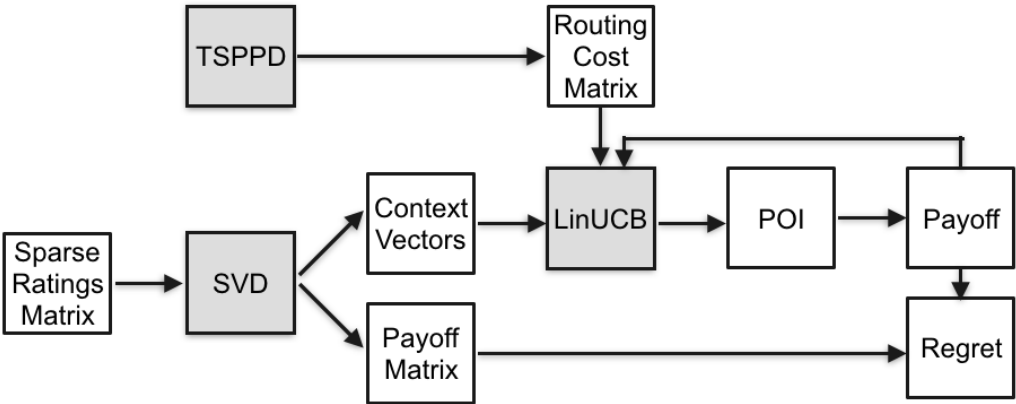


Figure 8: Recommendation engine framework.

We use a UCB contextual bandit algorithm (Li et al., 2010, 2017) modified to include routing cost in the feature vector shown in Algorithm 1. A new trial is initiated when a traveler accessing an MOD service from a pickup location requests recommendations for a destination, or when certain incident conditions warrant having the passenger change their original destination to a new location. Algorithm 1 calculates the upper confidence bound (UCB) for each destination and chooses the one with the highest UCB to recommend to the traveler.

Algorithm 1. UCB-GLM with routing costs	
Input:	$\alpha \in \mathbb{R}_+$ (for $\delta = 0.05$ (95% upper bound), $\alpha = 2.358$), minimum initial random samples τ
0:	Initialize: for first $t \leq \tau$ samples randomly choose $a_t \in \mathcal{A}_t$ and set $V_{\tau+1} = \sum_{t=1}^{\tau} X_t X_t'$, where $X_t = [1, x_{t1}, \Delta_t]$ includes the chosen destination's rating (x_{t1}) and increase in routing cost Δ_t for user t
for $t = \tau + 1, \tau + 2, \dots, T$ do	
1:	$\hat{\theta}_t \leftarrow \sum_{i=1}^{t-1} (Y_i - \mu(X_i' \theta)) X_i = 0$, where $\mu(X_i' \theta) = X_i' \theta$, if linear $\mu(X_i' \theta) = (1 + \exp(-X_i' \theta))^{-1}$, if logistic
2:	$a_t = \operatorname{argmax}_a \left\{ \mu(X_{t,a}' \hat{\theta}_t) + \alpha \sqrt{X_{t,a}' V_t^{-1} X_{t,a}} \right\}$
3:	Observe Y_t , let $X_t \leftarrow X_{t,a_t}, V_{t+1} \leftarrow V_t + X_t X_t'$
end for	

We modify the UCB measure to include a route cost increase component, where Δ_{ta} is measured as the increase in travel cost of adding the candidate destination to an existing route that is being served by a vehicle. The effective feature vector becomes $[x_{ta_t}, \Delta_{ta}]$. The variable p_{ta} is the expected payoff, $A_a \equiv D_a^T D_a + I_d$, where D_a is a design matrix of dimension $m \times (d + 1)$ corresponding to m training inputs, and b_a is the corresponding response vector (e.g. acceptance/rejection feedback). The parameter δ is a reliability measure used to tune the algorithm, where $\alpha = 1 + \sqrt{\ln\left(\frac{2}{\delta}\right)}/2$. The parameter $\hat{\theta}_a$ is the vector of estimated linear coefficients corresponding to the attributes and routing cost increase Δ_{ta} .

The algorithm has several properties that make it suitable for this experiment. First, it is computationally efficient as the number of arms is always linear and the number of features is at most cubic.

Additionally, the algorithm performs well on a robust arm set and remains competent if A_α is not too large (Li et al., 2010).

For evaluation of the algorithm performance, an offline analysis of its simulated performance will be used. A typical metric of offline evaluation of regret is shown in Eq. (2).

$$\frac{R_T}{T} = \frac{\sum_{t=1}^T \left(\mu \left(X'_{t,a_t} \theta^* \right) - \mu \left(X'_{t,a_t} \theta^* \right) \right)}{T} \quad (2)$$

R_T denotes the regret after T trials.

An illustration of this is shown for Yelp restaurants in Las Vegas shown in Figure 9. A database of 4029 establishments and user data covering approximately 1% of the fields was used as input to SVD. The resulting features matrices based on ratings between 1 and 5 were used to initialize a simulation drawing randomly from 10 individuals in choosing destinations over 100 trials. An illustration of the average regret trajectory realized by the trials (relative to the SVD payoffs), without any consideration of routing costs, is shown in Figure 10 for three different values of δ . The algorithm leading to the lower cumulative regret (smallest area under trajectory) would be the best performing algorithm.

Obviously, having such a dynamic component as the route cost increase is expected to dampen the learning rate of any learning algorithm. But by how much? Which variables contribute more to it? Before we design experiments to better understand its effect, we next introduce the routing algorithm used.




Figure 9: Heatmap of establishments with Yelp Open Data in Las Vegas.





Figure 10: Average regret trajectory without routing cost consideration.

4.2. Routing subproblem

The variable $\Delta_{t\alpha}$ is determined by considering the additional routing cost of the MOD service to take the customer to destination α , assuming that the routing cost increase converts to a fare price for accessing the destination for that customer. This is not a trivial task, as illustrated in Figure 11. In this example, a vehicle with passenger capacity of 2 has a purple passenger on-board with a blue passenger queued up when the red passenger makes a request for a destination recommendation. Without the red passenger, the original optimal route would be the length of the dashed black lines. If location 1 is recommended, then the optimal route would be the dashed red line. If location 2 is recommended, the optimal route would be the dotted red line. It is apparent that different locations can lead to drastically different sequences.



Suppose there are 3 passengers generated in a trial in a region:

Node numbers	Origin zone	Destination zone
0 (initial vehicle loc.)	67	
1	14	64
2	58	99

3	5	92
---	---	----

Using insertion heuristic in ascending order of nodes (pickups and drop-offs are examined simultaneously as a “double insertion”) results in the following sequence of node numbers:

(0,2,1,3,4,6,5)

The vehicle’s location (maintaining having all 3 customers) is distance-distributed between (0,2) (4.2 min), (2,1) (17.0 min), (1,3) (4.2 min), (3,4) (18.2 min). A random seed of 36.5% leads to the vehicle being on (2,1), and another random value of 65.6% generates the exact vehicle location at zone 58. We have the following trial input for a trial for the contextual bandit:

Vehicle location zone: 58

Passengers on-board: {2}

Remaining destination sequence: (1P, 3P, 1D, 3D, 2D)

Remaining route length: 71.3 minutes

If a new customer is generated at zone 28, each zone in this study area would be evaluated as one of 100 arms to determine which to recommend to the user.




Figure 12: Insertion heuristic solution.

Insertion heuristics are common heuristic used in many dynamic routing problems (e.g. Gendreau et al., 1992; Berbeglia et al., 2010; Jung and Jayakrishnan, 2014; Chow and Liu, 2012).

5. Computational experiments

5.1. Experimental setup

Having implemented a prototype of the recommender system, we demonstrate how it fares when considering routing constraints compared to ignoring it even though it is there. For these tests, a simple experimental design is considered.

In most real-life applications, we have access to information (context) that can be used to make a better decision. In the setting of a contextual bandit, we can think of features of each sample as the context and rewards are either 1 or 0 depending on whether we predict the class label correctly or not (i.e. whether the agent chooses to visit the system recommended restaurant or not).

Data for contextual bandits is not easily available and is hard to work with in terms of evaluation, but it is possible to turn any multi-label classification dataset into contextual bandits by revealing only the class for one label per observation. Doing this also has the advantage that we know also the rewards for all other arms when we want to evaluate a policy, but we can choose not to reveal them to the agent to simulate a real scenario.

The experiments here consist of iterating over a multi-label dataset, letting each policy make choices as it passes through the data and observes rewards for the actions it chooses, and recording the rewards that they receive. The base classification algorithm used is logistic regression.

The simulated dataset contains people's choices on whether to visit or not one of the three restaurants (the goal being to learn to suggest restaurants based on user previous history, i.e. dataset contains information on whether user visited_chinese, visited_american restaurants).

Given a set of predefined zones N , travel times $t_{ij}, i, j \in N$, a set of existing customers Q that are being served by the vehicle, where $o_i, i \in Q$, indicates zone of origin, $d_i, i \in Q$, indicates zone of destination, and $b_i \in \{0,1\}, i \in Q$, where $b_i = 1$ means that passenger i has already been picked up, and v_0 is the vehicle's closest zone.

In this problem the agent needs to make a sequence of decisions in time $t = 1, 2, \dots, T$. At each time t the agent is given a set of K arms (restaurants), and it has to decide which arm to pull. After pulling an arm, it receives a reward of that arm and the rewards of other arms in that trial are unknown. Let $T = 10,000$ and $K = 3$. We can run Algorithm 1 with and without routing constraints on the simulated dataset. We only keep an observation when the bandit agrees on the arm choice of the randomized arm choices specified in the initial dataset.

For a trial scenario:

- 100 establishments uniformly distributed over a 10x10 square grid
- Distances measured by Euclidean distance
- Service is a single shuttle with passenger capacity
- Each trial is represented by a number of pre-assigned passengers (Poisson distribution, max 10) with uniformly random centroid OD locations, and uniformly random location along this tour; an initial vehicle location uniformly distributed over centroids; e.g.
- Customers are recommended one of the 100 zones as a destination. Each zone has a ranking between [3,5] (5 out of 5 is best; we assume that only locations with 3 or higher might be recommended).
- Each customer is assumed to be governed by a random utility model as follows:

$$U_{dn} = b_0 + b_1x_{1d} + b_2x_{2dn} + \varepsilon_{dn}$$

where

U_{dn} is the utility gained by user n accepting destination n (versus rejecting it)

x_{1d} is the destination d rating (3 to 5)

x_{2dn} is the routing cost increase for the shuttle to transport user n to destination d (min)

ε_{dn} is a Gumbel-distributed error representing the unobservable utility of each person

For the experiment, we designate the parameters but assume they are unknown: $b_0 = -5, b_1 = 2, b_2 = -0.1$

For example, if a destination has a rating of 4 and a route cost increase of 15 minutes, the true probability of acceptance is assumed to be 81.8%. When destination rating drops 3 to, the acceptance probability becomes 37.8%. With rating of 4 and travel time increase of 30 minutes instead, the probability is 50%.

For the scenario, we use $\lambda = 1$, passenger capacity 4, travel time conversion of 3 (30 minute along one edge of grid), $\delta = 0.05$.

Ratings are randomly generated for the region as shown in Figure 13.

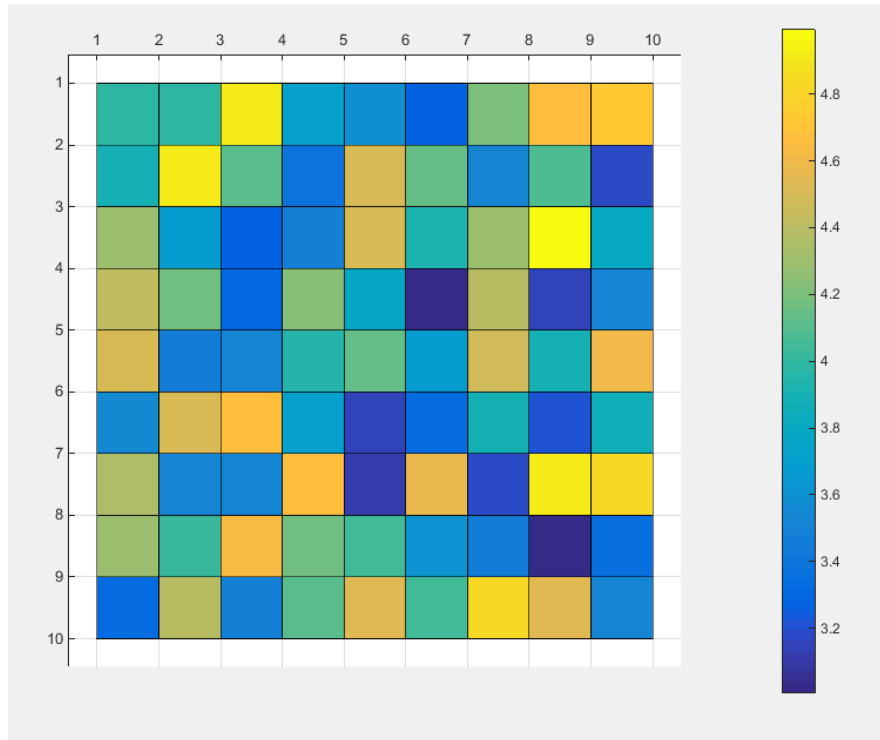


Figure 13: Simulated ratings (out of 5) across the 100 zones.

For the computational experiments, the prototype code can be found at <https://github.com/BUILTNYU/recommender-system>.

5.2. Results

We run the scenario for up to 2000 trials using a minimum sample of 12 to generate the initial estimate of $\hat{\theta}$. The estimated parameters after 2000 trials are $\hat{\theta} = (-5.913, 2.127, -0.088)$. This indicates that the algorithm is able to learn close to the true values of $\theta^* = (-5, 2, -0.1)$. We can also see how the distribution of recommendations is quite diverse over these trials as shown in Figure 14.

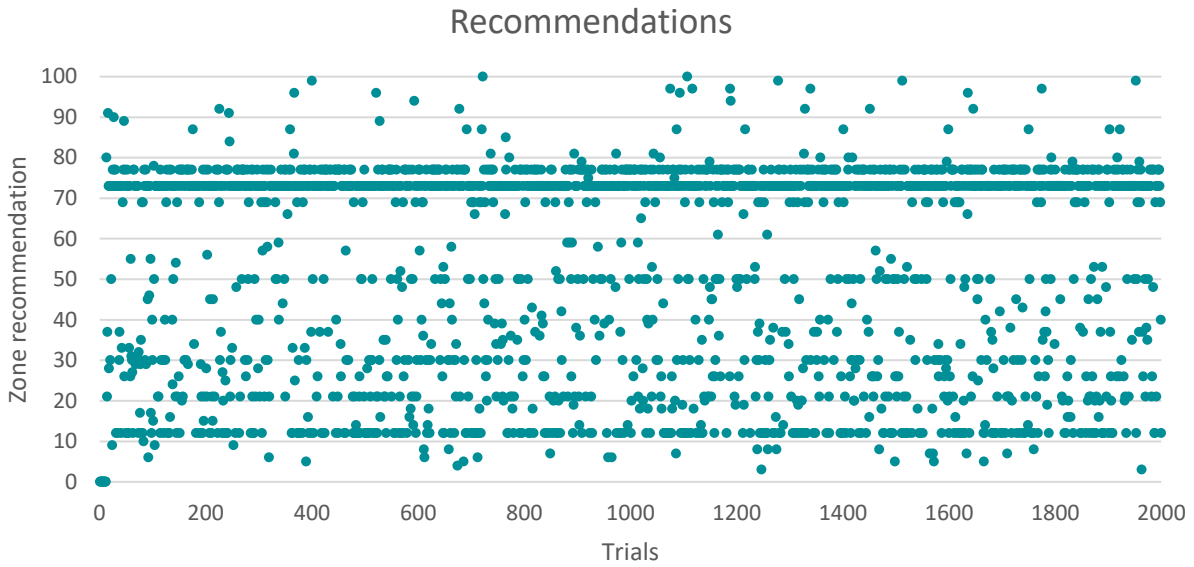


Figure 14: Realized recommendations for base scenario.

When we plot the spatial heat map of the most recommended locations, we get the following in Figure 15.

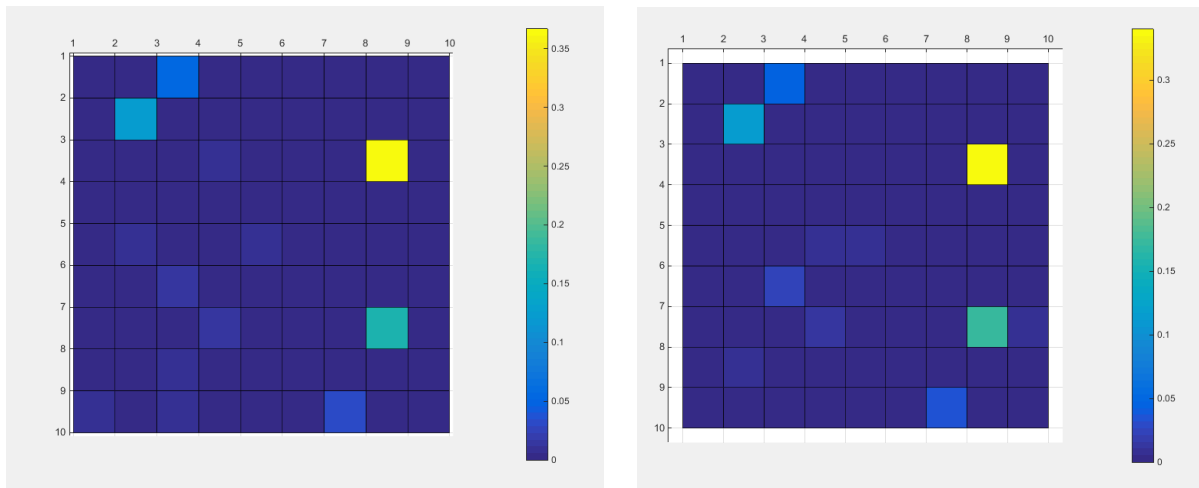


Figure 15: Distribution of recommended zones after (left) 500 trials, (right) 2000 trials.

As the figure shows, the recommendations become fairly stable when relying only on routing cost per individual and location rating for the destinations. These also generally match the trendy spots shown in Figure 13 (the ones in bright yellow). This suggests the effectiveness of the algorithm in recommending highly rated locations while accounting for routing costs.

The average regret is also computed for this case. The measure is compared against an alternative scenario where travel costs are increased by 4/3 over the base case.

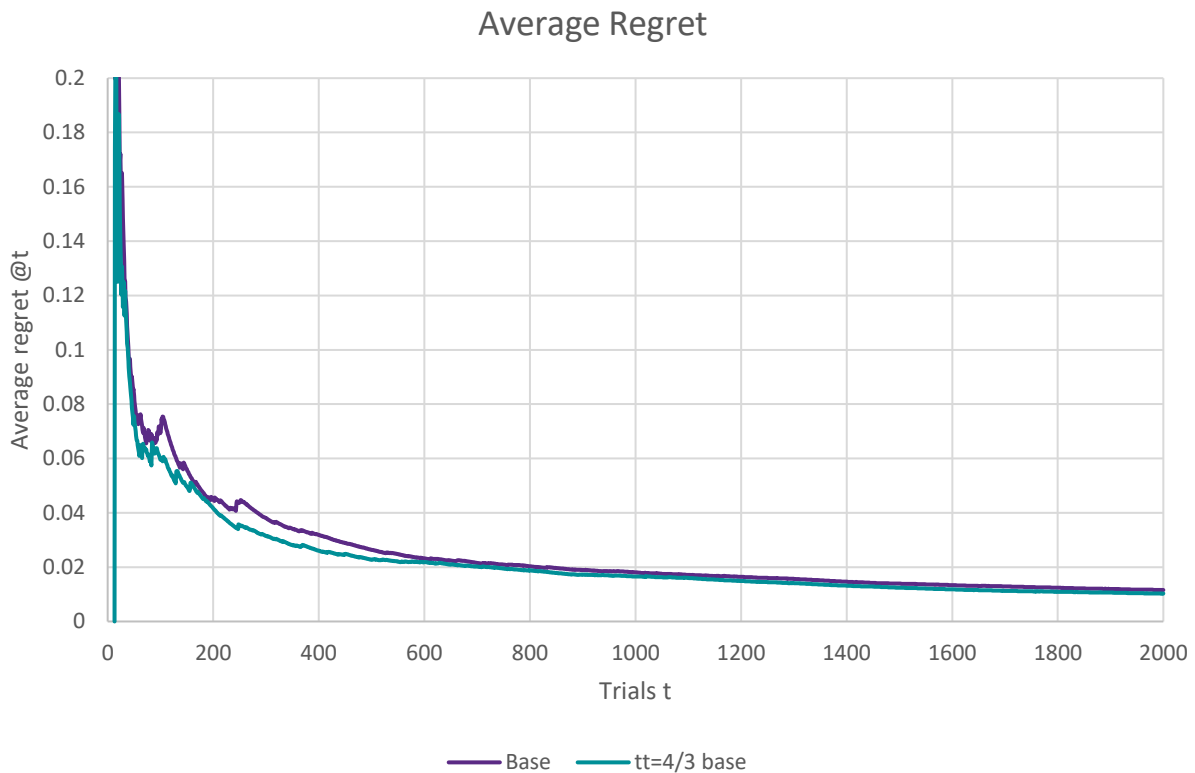


Figure 16: Average regret comparison.

We can see that the regret performance improves slightly when travel times are increased. This is interesting because Figure 17 suggests that there is more spatial distribution occurring. The increased spatial distribution occurs because the travel costs end up higher and so the system would recommend locations closer to the user. This results in more circumstantial results. On the other hand, the improved regret performance suggests that the high travel costs lead to less exploration by the algorithm needed to hone in on the best destinations. This is promising because it suggests that the benefit of using such a recommender system increases when the need (with increased travel costs) is greater.




Figure 17: Spatial distribution of recommendations when $tt = 4/3$ base.

6. City-wide deployment, deliverables, and technology transfer

6.1. Deployment

In this section some guidelines are provided to readers interested in deploying this system for pilot MOD services to incorporate destination recommender systems. There are many recommendation systems available for problems like shopping, online video entertainment, games etc. Restaurants & Dining is one area where there is a big opportunity to recommend dining options to users based on their preferences as well as historical data.

Yelp is a massive platform for crowdsourcing reviews of businesses such as restaurants and bars. Users of Yelp engage and interact with the application through searching businesses, writing reviews, rating businesses, connecting with other users, and “checking in” at businesses. The Yelp (2017) dataset has more information among the users, reviews and businesses. The Yelp dataset contains not only restaurant reviews, but also user-level information on their preferred restaurants.

Using Yelp’s dataset, we can extract collaborative and content-based features to identify customer and restaurant profiles. Based on the Yelp data, we can run the system with user observations and recommendations over time.

The dataset contains five different tables: User, Business, Review, Check-In and Tips. This information contains actual business, user, and users’ review data from >700 cities as JSON files. In total there are 156,000 businesses, 1,100,000 users, and 4,700,000 reviews. The ratings users give to businesses range from 1-5 as discrete values as a number of review stars. This dataset spans more than 10 years of Yelp reviews. Example features are shown in Table 5.

Table 3: Raw Features from Yelp

Business Data Features	User Data Features	Review Data Features
<pre>{' type': 'business', ' business_id': (encrypted business id), ' name': (business name), ' neighborhoods': [(hood names)], ' full_address': (localized address), ' city': (city), ' state': (state), ' latitude': latitude, ' longitude': longitude, ' stars': (star rating, rounded to half-stars), ' review_count': review count, ' categories': [(localized category names)] ' open': True / False (corresponds to closed, not business hours), ' hours': { (day_of_week): { ' open': (HH:MM), ' close': (HH:MM) }, }, ' attributes': { (attribute_name): (attribute_value), }, }</pre>	<pre>{' type': 'user', ' user_id': (encrypted user id), ' name': (first name), ' review_count': (review count), ' average_stars': (floating point average, like 4.31), ' votes': {(vote type): (count)}, ' friends': [(friend user_ids)], ' elite': [(years_elite)], ' yelping_since': (date, formatted like '2012-03'), ' compliments': { (compliment_type): (num_compliments_of_this_typ e), ... }, ' fans': (num_fans), }</pre>	<pre>{' type': 'review', ' business_id': (encrypted business id), ' user_id': (encrypted user id), ' stars': (star rating, rounded to half-stars), ' text': (review text), ' date': (date, formatted like '2012-03-14'), ' votes': {(vote type): (count)}, }</pre>

6.2. Deliverables

The following products were developed from this project.

- Survey results conducted from the joint study with the UTEP group has been published in a conference paper (Cheu et al., 2018) that will be presented in Kansas City, MO
- A prototype recommender system with routing subproblem was implemented and is publicly available: <https://github.com/BUILTNYU/recommender-system>

6.3. Technology transfer

A number of tech transfer activities were conducted during the course of this project.

- Presented at NSF RCN Workshop organized by Anil Yazici at Stony Brook, topic on “Smart and Connected Communities and Aging Population” on April 20th . (April 19-20, 2018)
- Presented the developed web application during Research Expo 2018 at NYU Tandon. (April 27th, 2018), <http://engineering.nyu.edu/events/2018/04/27/2018-nyu-tdandon-school-engineering-research-expo>
- Presented at the Transit Techies NYC Meetup hosted by Sidewalk Labs
- Presented at the 15th International Conference on Travel Behavior Research at Santa Barbara, CA (July 15-20, 2018)

7. Conclusion

In this project we set out to better understand how contextual bandit algorithms can be used to provide destination recommendations for MOD services. One of the primary target audiences for this work is the senior population, as the population is aging and many mobility options are more difficult for seniors to access with limited information. Because of this need, we set out on three objectives.

The first objective of the project was to collect data in NYC as part of a joint effort with UTEP to better understand how the mobility preferences among the elderly scale from one city to another. Particularly in the context of emerging mobility services, there has been limited information thus far. Our survey provided a better understanding of these needs and preferences: in cities the elderly tend to more frequently use smartphone devices and uncertainties associated with a trip are the top concerns when traveling (cost, weather, on-time reliability). Accessibility is also a major challenge that needs to be overcome. These findings provide motivation for improving the intelligence of MOD services to better cater to the elderly among other population groups.

The second objective was to implement a prototype recommender system that can be adopted by MOD services. A code was developed based on the LinUCB algorithm from Li et al. (2010) with a modification to the payoff function to include increase in routing cost as an additional variable. The code is available on our public repository.

The third objective was to conduct some computational tests with the prototype system to evaluate our hypothesis that it is important to explicitly incorporate routing constraints and that these constraints will tend to worsen the performance of the algorithm. Despite running into data security issues during the project (as reported in our final quarterly progress report), we were able to conduct some preliminary assessments that suggest there is an impact.

We can go along several directions for future research:

- The computational tests can be conducted more rigorously to evaluate the effect of an underlying graph on the impact of the worst-case regret bound. This should allow us to better understand what elements impact the learning process the most, for example, whether it's density, clustering of origin-destination patterns, or fleet operating policies.
- We can conduct a pilot with a MOD service to build up the recommender system with real data. This would be most effective as the contextual bandit relies heavily on estimation of the coefficients, which depend on real data.

- More efficient routing algorithms can be embedded to see how that affects the learning algorithm; for example, there are tabu search, GRASP algorithms, genetic algorithms (e.g. Chow and Liu, 2012; Allahviranloo et al., 2014; Chow, 2014), and adaptive large neighborhood search heuristics.

References

- AARP, 2015. AARP Public Policy Institute, "How the Travel Patterns of Older Adults are Changing: Highlights from the 2009 National Household Travel Survey," *Fact Sheet*, 2015. [Online]. Available: <http://www.aarp.org>.
- Abbasi-Yadkori, Y., Pál, D., & Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems* (pp. 2312-2320).
- Adomavicius, G. and Tuzhilin, A., 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734-749.
- Agarwal, D., Chen, B.C., Elango, P., Motgi, N., Park, S.T., Ramakrishnan, R., Roy, S. and Zachariah, J., 2009. Online models for content optimization. In *Advances in Neural Information Processing Systems* (pp. 17-24).
- Allahviranloo, M., Chow, J. Y., & Recker, W. W. (2014). Selective vehicle routing problems under uncertainty without recourse. *Transportation Research Part E: Logistics and Transportation Review*, 62, 68-88.
- Alsnih, R., & Hensher, D. A. (2003). The mobility and accessibility expectations of seniors in an aging population. *Transportation Research Part A: Policy and Practice*, 37(10), 903-916.
- Arentze, T. and Timmermans, H., 2003. Modeling learning and adaptation processes in activity-travel choice A framework and numerical experiment. *Transportation* 30(1), 37-62.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov), 397-422.
- Bao, J., Zheng, Y., Wilkie, D., & Mokbel, M. (2015). Recommendations in location-based social networks: a survey. *Geoinformatica*, 19(3), 525-565.
- Berbeglia, G., Cordeau, J. F., & Laporte, G. (2010). Dynamic pickup and delivery problems. *European journal of operational research*, 202(1), 8-15.
- Berenguer, A., Goncalves, J., Hosio, S., Ferreira, D., Anagnostopoulos, T., & Kostakos, V. (2017). Are Smartphones Ubiquitous?: An in-depth survey of smartphone adoption by seniors. *IEEE Consumer Electronics Magazine*, 6(1), 104-110.

- Bird, D. C., Freund, K., Fortinsky, R. H., Staplin, L., West, B. A., Bergen, G., & Downs, J. (2017). Driving self-regulation and ride service utilization in a multicomunity, multistate sample of US older adults. *Traffic Injury Prevention, 18*(3), 267-272.
- Brunato, M. and Battiti, R., 2003. PILGRIM: A Location Broker and Mobility-Aware Recommendation System. In proceedings of IEEE PerComp2003 (March 23-26, 2003), 265-272.
- Cheu, R.L., Villanueva-Rosales, N., Nunez-Mchiri, G.G., Vechione, M., Vargas-Acosta, R.A., Marrufo, C., Jimenez-Velasco, M.G., Gurbuz, O., Dmitriyeva, A., Chow, J.Y.J., 2018. Bringing smart mobility services to seniors: challenges and solutions in El Paso, TX, and New York, NY, In: Proc. IEEE ISC2, Kansas City, MO.
- Chow, J. Y. J. (2014). Activity-based travel scenario analysis with routing problem reoptimization. *Computer-Aided Civil and Infrastructure Engineering, 29*(2), 91-106.
- Chow, J.Y.J., and Liu, H., 2012. Generalized profitable tour problems for online activity routing system. *Transportation Research Record 2284*, 1-9.
- Chu, W., Li, L., Reyzin, L., & Schapire, R. (2011, June). Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (pp. 208-214).
- Crains, 2016. Access-a-ride wastes up to \$100 million a year, experts say. Crain's New York, Sept. 20, 2016.
- Deloitte, 2015. Smart mobility: reducing congestion and fostering faster, greener, and cheaper transportation options. Deloitte Smart Mobility Research Report, <https://dupress.deloitte.com/dup-us-en/industry/public-sector/smart-mobility-trends.html>, last accessed Oct. 8, 2017.
- Filippi, S., Cappe, O., Garivier, A., & Szepesvári, C. (2010). Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems* (pp. 586-594).
- Freund, K., & McKnight, A. J. (1997). Independent transportation network: alternative transportation for the elderly. TRANSIT-IDEA Program Project Final Report.
- Geerlings, M. I., Jonker, C., Bouter, L. M., Adèr, H. J., & Schmand, B. (1999). Association between memory complaints and incident Alzheimer's disease in elderly people with normal baseline cognition. *American Journal of Psychiatry, 156*(4), 531-537.
- Gendreau, M., Hertz, A., & Laporte, G. (1992). New insertion and postoptimization procedures for the traveling salesman problem. *Operations Research, 40*(6), 1086-1094.

Hanley, I. G. (1981). The use of signposts and active training to modify ward disorientation in elderly patients. *Journal of Behavior Therapy and Experimental Psychiatry*, 12(3), 241-247.

Jun, K. S., Bhargava, A., Nowak, R., & Willett, R. (2017). Scalable Generalized Linear Bandits: Online Computation and Hashing. In *Advances in Neural Information Processing Systems* (pp. 99-109).

Jung, J., & Jayakrishnan, R. (2014). Simulation framework for modeling large-scale flexible transit systems. *Transportation Research Record: Journal of the Transportation Research Board*, (2466), 31-41.

Li, L., Chu, W., Langford, J. and Schapire, R.E., 2010, April. A contextual-bandit approach to personalized news article recommendation. In Proceedings of the 19th international conference on World wide web (pp. 661-670). ACM.

Li, L., Lu, Y., & Zhou, D. (2017). Provably optimal algorithms for generalized linear contextual bandits. *arXiv preprint arXiv:1703.00048*.

March, L. M., Brnabic, A. J., Skinner, J. C., Schwarz, J. M., Finnegan, T., Druce, J., & Brooks, P. M. (1998). Musculoskeletal disability among elderly people in the community. *Medical Journal of Australia*, 168(9), 439-442.

NMHC (2015). State Distribution of Apartment Residents, <https://www.nmhc.org/research-insight/quick-facts-figures/quick-facts-resident-demographics/>

NYCHA (2012). Facilities and Service Centers, <https://data.cityofnewyork.us/Housing-Development/NYCHA-Facilities-and-Service-Centers/d4iy-9uh7/data>

Phithakkitnukoon, S., Veloso, M., Bento, C., Biderman, A. and Ratti, C., 2010, November. Taxi-Aware Map: Identifying and Predicting Vacant Taxis in the City. In Aml (pp. 86-95).

Rehman, F., Khalid, O., & Madani, S. A. (2017). A comparative study of location-based recommendation systems. *The Knowledge Engineering Review*, 32.

Ritter, A. S., Straight, A., & Evans, E. L. (2002). Understanding senior transportation: Report and analysis of a survey of consumers age 50. AARP, Public Policy Institute.

Rusmevichientong, P., & Tsitsiklis, J. N. (2010). Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2), 395-411.

Sayarshad, H. R., & Chow, J. Y. J. (2015). A scalable non-myopic dynamic dial-a-ride and pricing problem. *Transportation Research Part B: Methodological*, 81, 539-554.

SCOPE (2016). Sarasota County Senior Transportation Needs Assessment: Survey Results and Analysis.

Silvis, J., & Niemeier, D. (2009). Social network and dwelling characteristics that influence ridesharing behavior of seniors. *Transportation Research Record* (2118), 47-54.

United States Census Bureau, "Facts for Features: Older Americans Month: May 2017," *United States Census Bureau*, 2017. [Online]. Available: <https://www.census.gov/newsroom/facts-for-features/2017/cb17-ff08.html>

Vodopivec, N., Tobias, D., Miller-Hooks, E., Schonfeld, P. and Mohebbi, M., 2015. Taxis as a recourse option for ridesharing services. *Transportation Research Record* 2536, 86-97.

Yelp, 2017. Yelp Open Dataset, <https://www.yelp.com/dataset>, last accessed Oct. 8, 2017.