

Learning-Based Travel Prediction in Urban Road Network Resilience Optimization

Guancheng Qiu¹, Amrita Gupta², Caleb Robinson², Shuo Feng¹, Bistra Dilkina¹

¹ University of Southern California

² Georgia Institute of Technology

gqiumath@gmail.com, agupta375@gatech.edu, drobinson67@gatech.edu, shuo.s.feng@gmail.com, dilkina@usc.edu

Abstract

Urban mobility is a key part of routine operations in cities, but is increasingly at risk due to floods. To mitigate these risks, urban planning and disaster management agencies must anticipate the potential impacts of road network connectivity losses on travel flows between different locations in an urban area. However, detailed travel flow data are not widely available, necessitating the use of models for estimating origin-destination travel volumes based on geospatial and socioeconomic features. This paper examines several travel demand prediction models in the context of their suitability for informing road network resilience planning. We first evaluate the capacity of these models to capture census-tract-level urban travel demand patterns. We then use the predicted travel flows as input to a budget-constrained optimization scheme for identifying which roads to reinforce against flooding in order to preserve connectivity. We find that the road upgrade prioritization performs equally well when provided the ground truth travel demand data and moderately accurate predictions from the learning-based models. Thus, models with moderately lower prediction accuracy but with other computational and practical advantages may be favored for specific decision processes.

Introduction

The movement of individuals between urbanized areas accompanies nearly all aspects of daily life. Reliable information about travel demand is crucial to ensuring and improving the performance of transportation infrastructure. There has been a great deal of research on travel flow prediction, or estimating the number of trips taken between pairs of locations (Liu et al. 2020) given demographic, socioeconomic and/or geographic information about each location. Models developed for this purpose range from simple traditional ones with few parameters to complex ones capable of learning complex interactions from a large set of variables in order to more accurately capture the structure of mobility patterns (Spadon et al. 2019; Lenormand et al. 2012; Liu et al. 2020). The estimated origin-destination (OD) flows produced from such travel demand models can help decision-makers understand the use of road infrastructure and plan for its future.

One important planning problem is that of improving the resilience of urban mobility to disruptions due to extreme weather events such as floods. Direct flood damages in the US currently average \$9bn each year (National Weather Service), including critical public infrastructures such as highways, roads, bridges, and utilities. Such damage poses challenges to meeting travel demand associated with evacuation, emergency response, disaster relief distribution, and routine socioeconomic activity. For example, roads between the Mozambican capital Maputo and the rest of the country remained unusable for nearly a year after devastating floods in 2000, causing economic growth to come to a halt (Chinowsky and Arndt 2012). In order to avoid such severe outcomes, policy-makers must be able to assess how flooding threatens urban travel flows and plan where to allocate resources towards mitigation.

While there has been considerable prior work on predicting ground truth travel flows, there has been little to no assessment of how prediction errors might propagate and impact downstream decision making. Motivated by this important, real-world use case of travel flow prediction models, *we propose an evaluation pipeline for assessing a travel flow model's ability to guide infrastructure investments for disaster mitigation planning in urban areas.* We first compare traditional and learning-based models for predicting census-tract level travel flow volumes among census tracts in 3 urban areas in the U.S. We show that a new variant of a commonly-used trip distribution model, the gravity model, as well as a random forest (Breiman 2001) model using a rich set of landscape and socioeconomic features are more accurate than other models at predicting census-tract-level vehicular flows in our study areas. Next, we use the predicted travel flows as inputs to a greedy-based procedure for road network fortification planning, and assess the quality of the obtained plans with respect to the ground truth mobility flows. By examining whether the resulting recommended fortification plans are sensitive to errors in the predictions of the proposed travel demand models, we evaluate these models on the basis of their suitability for use in such a decision pipeline.

The proposed evaluation pipeline also makes a valuable contribution to spatial flood risk assessment by quantifying the *mobility* impacts of flooding. Several recent studies have measured the impacts of floods on transportation by

spatially intersecting flood hazard maps with road networks and reporting the total length of roadways within flood zones (Hankin et al. 2016; Kulp and Strauss 2017; Gupta, Robinson, and Dilkina 2018). However, this metric does not capture the degree of functional loss in the road network due to its exposure to the hazard (Pant, Hall, and Blainey 2016). The potential for functional losses in connectivity partly depends on the network’s topology: graph theoretic indicators such as edge or node centrality (Casali and Heinemann 2019) can be used to identify key road segments or junctions, and metrics like graph clustering coefficients can quantify overall network connectivity (Zhang and Alipour 2019). However, the magnitude of impact on mobility also depends on the flow along the edges exposed to hazard, which is addressed in this study.

Problem Formulation

We assume we are tasked with generating a budget allocation plan to protect urban mobility from flood hazards in a region in which ground truth OD travel volumes are unknown. We are given an undirected, uncapacitated graph $G = (V, E)$, where edges represent road segments, and vertices represent junctions or endpoints of the road segments. We are also given an OD (origin-destination) matrix T whose (i, j) -th entry contains the average number of daily trips from census tract i to tract j over the road network, which we refer to as the travel demand from tract i to j . Travel demand within a census tract is not taken into account.

With the OD flows predicted by a travel demand model, \tilde{T} , commuting flows from tract i to tract j are assumed to be evenly distributed among all possible paths from a road junction in tract i to one in tract j . We assume that as long as a path exists between a pair of road junctions, the travel demand between them in both directions is satisfied. We are given a cost function $c : E \rightarrow \mathbb{R}_+$ on the edges that reflects the projected cost of upgrading a given road segment to ensure it withstands a given flood scenario. The planner’s task is to decide which road segments to upgrade through the allocation of a budget B , such that the maximum travel demand can be satisfied.

Recent works have proposed different algorithmic solution approaches for this challenging mobility optimization problem, including a mixed-integer program formulation, a greedy heuristic algorithm, and a more complex heuristic algorithm based on supermodularity (Gupta, Robinson, and Dilkina 2018; Gupta and Dilkina 2019). However, a more fundamental challenge to solving this problem arises from the need to have access to the OD travel flow matrix T . Such travel demand data are rarely recorded, leading many research efforts to rely on cell phone data (Yang et al. 2014; Gupta, Robinson, and Dilkina 2018) to infer mobility patterns, or to use a subset of travel flows attributed to commuting trips based on census data. These data sources are also limited in availability and comprehensiveness, motivating the use of predictive models for travel demand T_{ij} between locations i and j at appropriate geographic scales for mobility resilience planning. In the following section we will

explore several models for estimating these flows.

Using the estimate of commuting flows on each edge of the road network as data for the optimization problem described above, we obtain a road infrastructure fortification plan tailored to protecting the predicted travel demand pattern. Specifically, we will use the simple greedy approach in (Gupta, Robinson, and Dilkina 2018) for minimizing the number of infeasible trips. We then assess the efficacy of the resulting plan by computing how much of the held-out ground truth travel demand remains feasible in the event of a flood hazard. In this work, we adopt the common practice of ignoring trips within a zone.

Related Work in Commuting Flow Prediction

Travel demand patterns over a study region can be summarized by an OD matrix T whose (i, j) -th entry contains the number of trips originating in zone i and ending in zone j within some time frame. Modeling these flows is often split into two parts: trip generation, which estimates the total number of trips leaving from or arriving at a given zone; and trip distribution, which characterizes what proportion of the trips generated for a given zone go to or come from each other zone.

Trip Distribution Models

Trip distribution models characterize the conditional probability $P(j|i)$ that a trip starting in zone i ends in zone j , based on features of the origin and destination zones and various assumptions about what other factors impact human mobility (Lenormand, Bassolas, and Ramasco 2016). An estimate of the number of trips from zone i to zone j , T_{ij} , is given by

$$\tilde{T}_{ij} = T_i P(j|i), \quad (1)$$

where T_i is the number of trips leaving zone i . T_i is often estimated by a production function $\tilde{T}_i = \lambda m_i$ where λ is a parameter that can be fitted. Historically, many models have been proposed for this task, but many of these fall into two major categories of approaches.

Gravity Models Gravity models (Carey 1867; Zipf 1946; Erlander and Stewart 1990) assume the probability P_{ij} that a trip begins in zone i and ends in zone j is proportional to the product of populations of the two zones, and inversely proportional to an exponential or power function of the distance d_{ij} between the zones, where d_{ij} can be the great-circle distance, Euclidean distance or travel distance between two zones. In this work, we consider variants using Euclidean distance and the shortest-path (in the given road network) travel distance between two census tracts’ centroids.

$$P_{ij} \propto \frac{m_i m_j}{e^{\beta d_{ij}}} \quad (2)$$

$$P_{ij} \propto \frac{m_i m_j}{d_{ij}^\beta} \quad (3)$$

where β is a parameter that can be adjusted, and P_{ij} is normalized so that $\sum_i \sum_j P_{ij} = 1$. We adopt the common practice of ignoring trips within a zone, i.e., for any i , $P_{ii} = 0$.

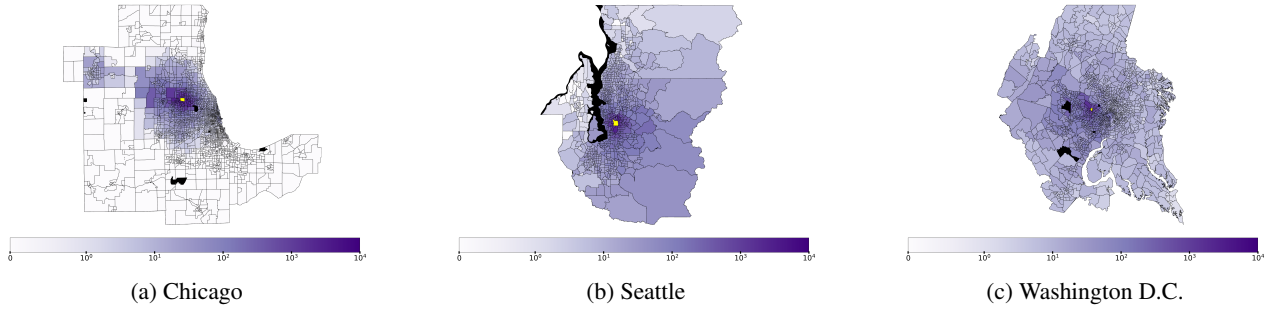


Figure 1: Number of outgoing trips from a single origin census tract (yellow) to every other census tract in each study area; black census tracts have no population.

Intervening Opportunities Models The second main family of human mobility models encompasses different variants of intervening opportunities models (Stouffer 1940). The number of intervening opportunities between zone i and zone j , s_{ij} , refers to the total number of jobs located closer to zone i than zone j is. We approximate this quantity with total number of jobs in all zones that are closer to zone i than zone j :

$$s_{ij} = \sum_{k: d_{ik} < d_{ij}} s_k \quad (4)$$

where s_k is the number of jobs in zone k . The number of jobs in a zone is sometimes estimated by the population of the zone, a convention we also adopt in this work.

Intervening opportunities trip distribution models approximate $P(j|i)$, the conditional probability that a trip leaving from zone i will go to zone j . Note that $P(j|i) = \frac{P_{ij}}{\sum_j P_{ij}}$. Different intervening opportunities models include:

- **Schneider’s intervening opportunities model** (Schneider 1959):

$$P(j|i) = e^{-\gamma s_{ij}} - e^{-\gamma(s_{ij} + m_j)}, \quad (5)$$

- **Radiation model** (Simini et al. 2012):

$$P(j|i) = \frac{m_i m_j}{(m_i + s_{ij})(m_i + m_j + s_{ij})}, \quad (6)$$

- **Extended Radiation model** (Yang et al. 2014):

$$P(j|i) = \frac{[(m_i + m_j + s_{ij})^\alpha - (m_i + s_{ij})^\alpha](m_i^\alpha + 1)}{[(m_i + s_{ij})^\alpha + 1][(m_i + m_j + s_{ij})^\alpha + 1]} \quad (7)$$

Learning-Based Travel Demand Models

One drawback of the traditional models is that the α , β , γ and λ parameters that are fit to reflect travel demand patterns in one region do not generalize well to other regions due to the rigid functional forms of the models and reliance on a small set of features. Here, we present two models that incorporate a wider set of features and capture more varied functional relationships between those features and the travel demand. Specifically, given data consisting of a set of

census tracts (zones), features F for each zone, joint features J between pairs of zones, and ground truth pairwise OD flows between zones over some time horizon, our goal is to learn a function $f(F_i, F_j, J_{ij}) = \tilde{T}_{ij}$ for predicting the OD flows for new areas in which we do not know the ground truth.

Extended Gravity Model

We consider a model that generalizes the gravity model to incorporate a much wider set of features (Anderson 2010). This extended gravity model adds power laws of additional features to the original gravity model, expressed as

$$\tilde{T}_{ij} = \beta \prod_l \phi_l^{\alpha_l}(i) \prod_m \phi_m^{\alpha_m}(j) \prod_n \phi_n^{\alpha_n}(i, j) f(d_{ij}) \quad (8)$$

where ϕ_l are features of the origin tract, ϕ_m are features of the destination tract, and ϕ_n are features related to both the origin and the destination (except the distance, which is included in the decay function $f(d)$). $f(d)$ can have either the power form or the exponential form. The bias β , and α coefficients are variables to be fit.

We used Poisson regression to fit the gravity model (Flowerdew and Aitkin 1982). We approximate T_{ij} as a Poisson-distributed random variable with rate parameter λ_{ij} , and is estimated as follows:

$$\begin{aligned} \tilde{T}_{ij} = \lambda_{ij} = & \exp(\ln \beta + \ln f(d_{ij}) \\ & + \sum_l \alpha_l \ln \phi_l(i) + \sum_m \alpha_m \ln \phi_m(j) + \sum_n \alpha_n \ln \phi_n(i, j)) \end{aligned} \quad (9)$$

where the parameters can be fitted by linear regression with the unit deviance of Poisson distribution as the loss function (Jørgensen 1992). 1 is added to all numeric values of a feature if the values contain 0 before the logarithm transformation. This method prevents several problems associated with the log-normal model, which would fit the parameters via linear regression on the log-transformed features. Problems of this approach which are avoided by the Poisson model include: that the antilogarithm of $\ln(\tilde{T}_{ij})$ is a biased estimator; the assumption that T_{ij} is log-normally distributed; the assumption that variance is the same for all T_{ij} ; and the impact to outcomes by the positive values added

Feature Category	Zone-based Feature	Origin-Destination Pair Feature
Population	Population	Intervening population
	Population density	
Geographic	Area	Euclidean distance
	Open space area	
	Low intensity development area	Shortest-path travel distance
	Medium intensity development area	
	High intensity development area	
	Forested area	
Work	Employed population	Intervening jobs
	Unemployed population	
	Average commute time	Intervening income
	Number of jobs	
	Per capita income	

Table 1: Feature set for the machine learning methods.

to zero values in the data for logarithmic transformation (Flowerdew and Aitkin 1982).

A simpler form of the extended gravity model was proposed by Lowry (Lowry 1966; Blanco 1963), which also considers employment status and income level as features which can be proportional or inversely proportional to travel demand. Earlier work by Alonso (Alonso 1976) also established a power-decay gravity model where the population values are replaced by power laws of population values.

Random Forest

We also use a random forest model, which is one of the most accurate predictors used in a recent study (Spadon et al. 2019) on predicting travel demand between cities based on socioeconomic features. Compared to the aforementioned models, random forest is a black-box model that is computationally more complex and more challenging to interpret.

Experiments

We evaluate the proposed travel demand prediction models in study areas around 3 major US cities susceptible to flooding: Washington D.C., Chicago and Seattle. For each city, we define the study area extent as the geographic area covered by the Traffic Analysis Zones (TAZs) used to obtain ground truth origin-destination travel demand for the region. The study area around Washington D.C. consists of the District of Columbia and parts of Maryland, West Virginia and Virginia; the study area around Chicago consists of parts of Illinois, Wisconsin and Indiana; and the study area around Seattle is wholly contained within Washington.

Datasets

Features We use geographic and socioeconomic features at the census-tract level to fit our travel demand models (Table 1). The census tract extents and indices are those used in the 2010 U.S. census, clipped to the TAZ-based study area for each city. Numerical features pertaining to population, land cover, employment, per capita income etc. for each census tract are obtained from the American Community Survey, the 2011 Environmental Summaries, and the 2015 Longitudinal Employer-Household Dynamics datasets on Social Explorer. For each pair of origin census tract i and destination census tract j , we have 13 zone-based features F_i relating to the origin, 13 zone-based features F_j relating to the destination, and 5 joint features J_{ij} .

Ground Truth Travel Demand TAZs are special spatial units used by transportation officials to track traffic-related data. We use publicly available TAZ shapefiles for Chicago, Seattle, and Washington D.C. (Chicago Metropolitan Agency for Planning; Puget Sound Regional Council; National Capital Region Transportation Planning Board), as well as travel volume data in the form of OD matrices containing the number of trips from each TAZ to every other TAZ in a given day. Since TAZs are a different geographical parcellation from census tracts, we resample the raw OD flow data between TAZs to estimate OD flows between census tracts (Figure 1) using the following formula:

$$T_{ij} = \sum_a \sum_b T_{ab} p(a, i) p(b, j) \quad (10)$$

where a, b are TAZs, i, j are census tracts, and $p(a, i)$ denotes the proportion of area of TAZ a that overlaps with census tract i , and $p(b, j)$ similarly. In other words, the trips to and from a TAZ are assumed to be evenly distributed across its area.

Road Networks We obtain graph representations of the road networks in our study areas from OpenStreetMap. Edges represent road segments and nodes represent intersections or dead ends, and the graphs are directed and may contain multi-edges and loops. We simplify each road network graph by transforming all directed edges to undirected ones and replacing multiple edges between two nodes with a single edge, resulting in a simple undirected graph.

Flood Hazard The Federal Emergency Management Agency provides shapefiles specifying flood zones in the US and their types (Federal Emergency Management Agency). Only two types of zones are considered in this study: zones that are marked with a 0.2% annual chance of flooding (likely to be flooded once every 500 years), and zones that have mandatory flood insurance purchase requirements that have more than 0.2% annual chance of flooding. The length of each road segment at risk of flooding is then calculated as the length of it that intersects with these flood zones.

Model Training

We compare several variants of traditional travel prediction models: gravity models using exponential and power functions of Euclidean and shortest-path travel distance respectively; Schneider’s intervening opportunities model; a radiation model; and an extended radiation model. For each of the learning-based models, we use the OD flows from 2 of the 3 study areas as a training set and test the trained models on the OD flows of the third held-out study area. This process is repeated in 3 folds, corresponding to all 3 combinations of pairs of study areas. For the learning-based models, we use cross-validation to select the model hyperparameters. In all cases, census tracts with 0 population are excluded due to their incompatibility with the gravity models. Only inter-zone trips are predicted: for any i , $\hat{T}_{ii} = 0$.

Prediction Performance

We use 4 evaluation metrics to measure the prediction performance of our travel demand models in terms of the agree-

ment between the off-diagonal entries of the ground truth and predicted origin-destination travel demand matrices (Table 2). We report the normalized root mean square error (*NRMSE*) and coefficient of determination (r^2) which are commonly used evaluation metrics for regression models. In addition, we use two variants of the common part of commuters metric (*CPC*) (Gargiulo et al. 2012; Lenormand et al. 2012) widely used in travel prediction:

$$CPC(T, \tilde{T}) = \frac{2 \sum_{i,j=1}^n \min(T_{ij}, \tilde{T}_{ij})}{\sum_{i,j=1}^n T_{ij} + \sum_{i,j=1}^n \tilde{T}_{ij}} \quad (11)$$

The greater the agreement between the predicted travel flow volumes and the ground truth values, the closer the *CPC* is to 1. (Lenormand, Bassolas, and Ramasco 2016) recently proposed the common part of commuters according to distance CPC_d , which measures how well a model predicts the distribution of travel distance, disregarding specific origins and destinations. If N_k is the number of trips with distance between $2k - 2$ and $2k$ km, and \tilde{N}_k is the corresponding prediction:

$$CPC_d(T, \tilde{T}) = \frac{\sum_{k=1}^{\infty} \min(N_k, \tilde{N}_k)}{\sum_{i,j=1}^n T_{ij}} \quad (12)$$

which equals to 1 if, for every distance bin, the ground truth and the prediction have the same number of trips within the range; it equals to 0 if every trip from the ground truth data is within a different distance bin than all predicted trips.

The intervening opportunities models generally had poorer performance for predicting census-tract-level travel flows around a city. Among the gravity models, those using an exponentially decaying function of distance outperformed those using a power decay function. The learning-based travel flow prediction models outperformed all of the traditional models in all evaluation metrics. Random forest was the best performing model overall in terms of CPC_d , *NRMSE* and r^2 , while the extended gravity model with exponential decay was the best according to the *CPC* metrics. Note that the extended gravity model generally is much less time- and space-demanding than random forest.

Feature Importance

To inspect how much each learning-based model relied on each of the features listed in Table 1, we performed permutation feature importance analysis (Breiman 2001; Fisher, Rudin, and Dominici 2018) using *CPC* as performance metric. The results are shown in Figure 2. The shortest-path distance between travel origin and destination is an important feature, and in particular it is more important than the Euclidean distance across models and study areas. The number of jobs in the origin and destination zones were also found to be stable features of similar importance to all 3 learning-based models across all study areas. We also observe that random forest, an ensemble method, was less sensitive to noise in any one feature.

Spatial Distribution of Errors

Figure 3 illustrates the prediction errors made by three of the models (one of the best performing traditional models -

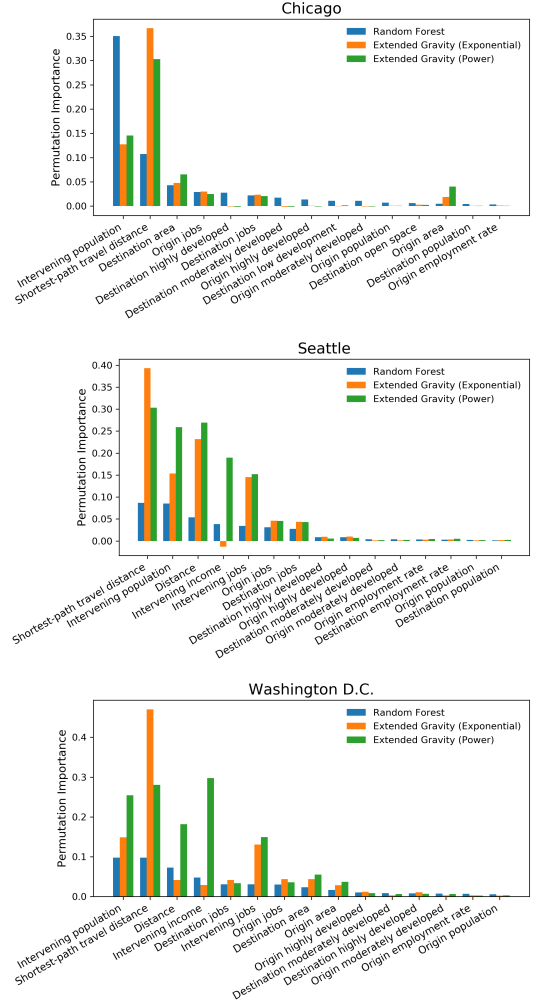


Figure 2: Permutation feature importance (PFI) for learning-based travel flow prediction models based on *CPC*. X axis is ordered by PFI of the random forest model.

Category	Method	CPC	CPC _d	NRMSE	r ²
Traditional models	Gravity, exponential decay, euclidean distance	0.583±0.032	0.860±0.063	6.099±2.854	0.338±0.118
	Gravity, exponential decay, travel distance	0.590±0.039	0.865±0.061	6.197±2.881	0.316±0.116
	Gravity, power decay, euclidean distance	0.552±0.070	0.790±0.081	6.993±3.711	0.155±0.257
	Gravity, power decay, travel distance	0.552±0.078	0.781±0.077	7.427±4.039	0.051±0.318
	Schneider's model	0.533±0.021	0.841±0.032	6.036±2.553	0.344±0.057
	Radiation model	0.297±0.046	0.430±0.036	17.660±7.619	-4.612±0.545
Learning based models	Extended radiation model	0.553±0.064	0.799±0.081	6.692±3.239	0.211±0.160
	Random Forest	0.654±0.068	0.907±0.084	5.287±2.529	0.506±0.095
	Extended gravity, exponential decay	0.658±0.067	0.879±0.134	5.597±3.260	0.462±0.221
	Extended gravity, power decay	0.629±0.060	0.826±0.160	6.216±3.606	0.340±0.269

Table 2: Average prediction performance. For the traditional trip distribution models, evaluation metrics are computed for a model fit to each of the 3 study areas and then averaged. For the learning-based models, the average is taken over the 3-fold tests. Higher is better for all metrics except for *NRMSE*.

gravity model with exponential decay, the extended gravity model with exponential decay, and the random forest model) on the number of outgoing trips from a single census tract in each study area. The gravity model tends to under-predict the travel demand to most destination census tracts, particularly to nearby census tracts. For example, the gravity model greatly underestimates the travel volume between the small origin tract in Washington D.C. and neighboring census tracts in the center of the study area. In contrast, the learning-based models are able to model this phenomenon, indeed over-correcting and leading to some overestimation of flows to some nearby tracts.

Road Network Optimization Decision Error

We next study the decision error incurred as a result of using the travel demand predictions of each of these models to inform a road upgrade optimization scheme. Since each of the study areas spans multiple counties, for the sake of scalability, one county from each study area is selected: Frederick, MD, Pierce, WA, and Lake, IN. These counties are selected since flooding is a significant threat in these areas according to the flood zone data. We extract the predicted origin-destination flows for census tracts within these counties from the predictions of their corresponding study areas.

Under a given flooding scenario, road segments that intersect with flood zones are considered to be unusable for travel, resulting in the originally connected road network graph becoming fragmented into multiple connected components and trips between nodes in different connected components becoming infeasible. The objective is to minimize the number of trips that may become infeasible this way by allocating a fixed budget towards reinforcing specific road segments against flooding, where the cost of reinforcing a road segment is assumed to be proportional to its length that lies within a flood zone. We employ a greedy algorithm to select which road segments to upgrade (Gupta, Robinson, and Dilkina 2018) in order of maximum benefit to cost ratio. Then, we assess the quality of each plan by using the ground truth OD flows to evaluate the number of remaining infeasible trips with the recommended upgrades in place. Commuting flows from tract i to tract j are assumed to be evenly distributed among all possible paths from a road junction in tract i to one in tract j .

Figure 4 shows the result of using OD flows predicted by the learning-based models, the gravity model with travel distances, and ground truth data to inform the decision process, and Figure 5 shows the spatial distribution of differences in restored trips. At the low budget levels shown, only a few critical road segments can be chosen for flood resilience upgrades and so segments that restore the most mobility must be correctly identified. The learning-based models perform comparably to each other but all outperform the simple gravity model in this regime, resulting in road upgrade plans that restore approximately 9000 additional trips at the 0.2% budget level, closely matching the quality of the plans obtained using the ground truth OD flows. These results give a strong indication that indeed predictive models for OD flows can be used in guiding urban transportation mitigation planning.

Conclusion

We find that including more relevant features and introducing flexibility into the functional form of travel demand prediction models can markedly improve travel flow prediction accuracy. Errors in predictions can affect subsequent decision processes differently. In this work, we consider a mobility resilience optimization problem that uses predicted travel demands as data. We observe that moderately accurate predictions by models such as the extended gravity model are sufficient for informing good decisions. Thus, in this scenario, the extended gravity model may be favored over the random forest model since it, despite performing moderately worse in some metrics than the latter, is far less computationally expensive. This shows that considering a specific decision process related to urban mobility can be important to assessing different models for predicting travel flows.

Acknowledgments

This work was supported by the U.S. Department of Homeland Security under Grant Award No. 2015-ST-061-CIRC01. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S Department of Homeland Security.

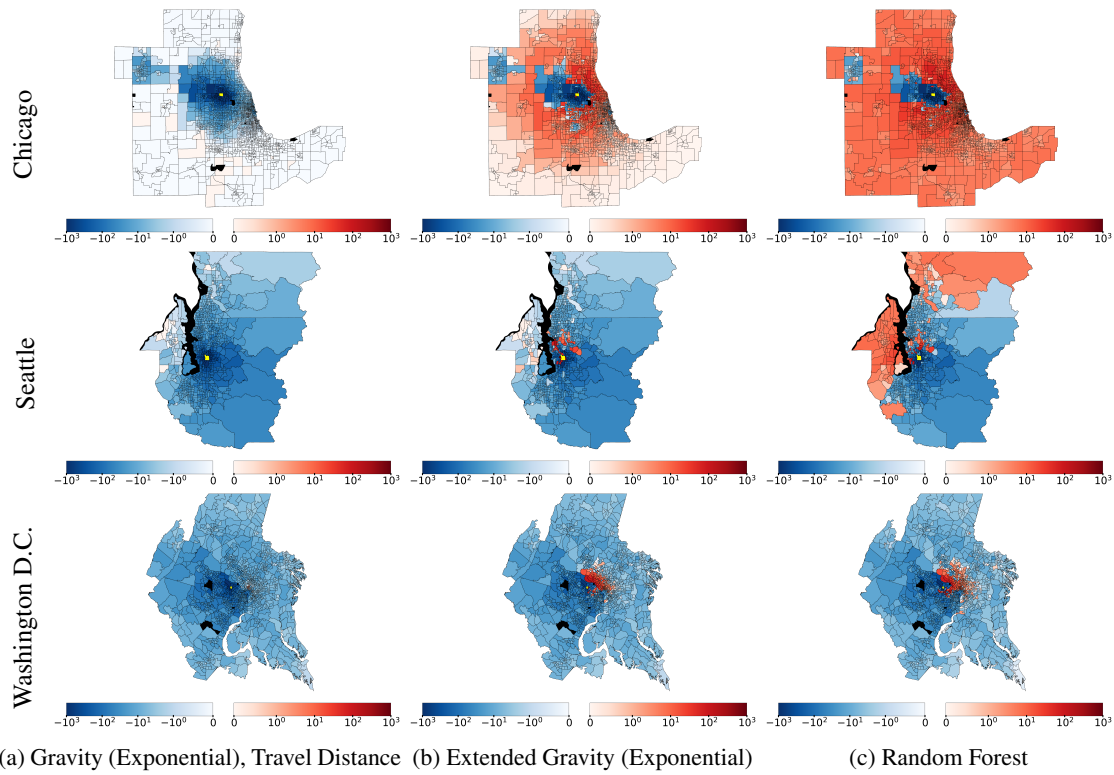


Figure 3: Prediction errors on number of outgoing trips from a single origin census tract (yellow) made by 3 different models.

References

- Alonso, W. 1976. *A theory of movements*. Inst. of Urban and Regional Devel.
- Anderson, J. E. 2010. The Gravity Model. Working Paper 16576, National Bureau of Economic Research. doi:10.3386/w16576. URL <http://www.nber.org/papers/w16576>.
- Blanco, C. 1963. The determinants of interstate population movements. *Journal of Regional Science* 5(1): 77–84.
- Breiman, L. 2001. Random forests. *Machine learning* 45(1): 5–32.
- Carey, H. C. 1867. *Principles of social science*, volume 3. JB Lippincott & Company.
- Casali, Y.; and Heinimann, H. R. 2019. A topological characterization of flooding impacts on the Zurich road network. *PLoS one* 14(7): e0220338.
- Chicago Metropolitan Agency for Planning. 2009. Modeling Zones. <https://datahub.cmap.illinois.gov/dataset/cmap-modeling-zone-systems>.
- Chinowsky, P.; and Arndt, C. 2012. Climate change and roads: A dynamic stressor–response model. *Review of Development Economics* 16(3): 448–462.
- Erlander, S.; and Stewart, N. F. 1990. *The gravity model in transportation analysis: theory and extensions*, volume 3. Vsp.
- Federal Emergency Management Agency. 2019. <https://msc.fema.gov/portal/advanceSearch>.




Figure 4: Number of infeasible trips in Pierce, WA after upgrades recommended by the greedy algorithm using different OD flow data at multiple budget levels.

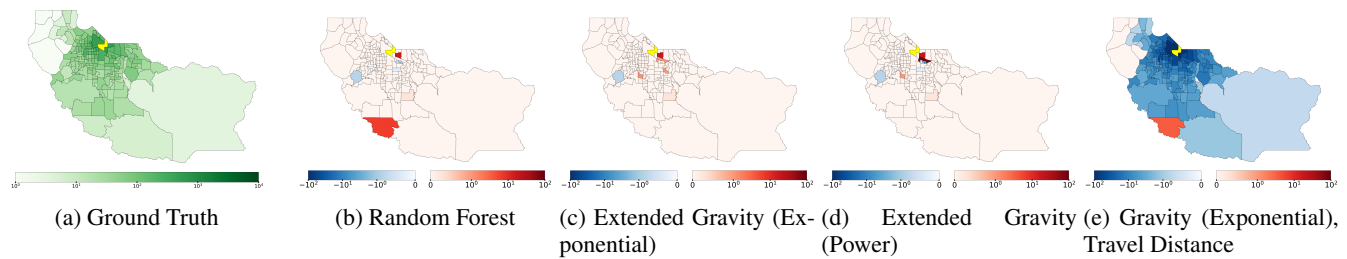


Figure 5: Number of restored infeasible trips from a single origin census tract (yellow) to every other census tract in Pierce, WA using ground truth flows (a), and differences in number of restored trips between fortification plans generated with predicted flows and that with the ground truth, given budget sufficient for upgrading 0.2% of the total length of flooded road segments.

Fisher, A.; Rudin, C.; and Dominici, F. 2018. Model class reliance: Variable importance measures for any machine learning model class, from the “Rashomon” perspective. *arXiv preprint arXiv:1801.01489* 68.

Flowerdew, R.; and Aitkin, M. 1982. A method of fitting the gravity model based on the Poisson distribution. *Journal of regional science* 22(2): 191–202.

Gargiulo, F.; Lenormand, M.; Huet, S.; and Espinosa, O. B. 2012. Commuting network models: Getting the essentials. *Journal of Artificial Societies & Social Simulation* 15(2): 6.

Gupta, A.; and Dilkina, B. 2019. Budget-Constrained Demand-Weighted Network Design for Resilient Infrastructure. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, 456–463. IEEE.

Gupta, A.; Robinson, C.; and Dilkina, B. 2018. Infrastructure resilience for climate adaptation. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 1–8.

Hankin, B.; Craigen, I.; Rogers, W.; Morphet, J.; Bailey, A.; and Whitehead, M. 2016. Flood Risk to the Strategic Road Network in England. In *E3S Web of Conferences*, volume 7, 10001. EDP Sciences.

Jørgensen, B. 1992. The theory of exponential dispersion models and analysis of deviance .

Kulp, S.; and Strauss, B. H. 2017. Rapid escalation of coastal flood exposure in US municipalities from sea level rise. *Climatic Change* 142(3-4): 477–489.

Lenormand, M.; Bassolas, A.; and Ramasco, J. J. 2016. Systematic comparison of trip distribution laws and models. *Journal of Transport Geography* 51: 158–169.

Lenormand, M.; Huet, S.; Gargiulo, F.; and Deffuant, G. 2012. A universal model of commuting networks. *PloS one* 7(10): e45985.

Liu, Z.; Miranda, F.; Xiong, W.; Yang, J.; Wang, Q.; and Silva, C. 2020. Learning geo-contextual embeddings for commuting flow prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 808–816.

Lowry, I. S. 1966. *Migration and metropolitan growth: two analytical models*. Chandler Pub. Co.

National Capital Region Transportation Planning Board. 2019. TAZ. <http://rtcdmwcog.opendata.arcgis.com/datasets/taz>.

National Weather Service. 2020. NWS Annual Flood Loss Summary Reports To U.S. Army Corps of Engineers. <https://www.weather.gov/water/>. Accessed: 2020-09-15.

Pant, R.; Hall, J. W.; and Blainey, S. P. 2016. Vulnerability assessment framework for interdependent critical infrastructures: case-study for Great Britain’s rail network. *European Journal of Transport and Infrastructure Research* 16(1).

Puget Sound Regional Council. 2010. Traffic Analysis Zone - 2010. <https://www.psrc.org/gis-shapefiles>.

Schneider, M. 1959. Gravity models and trip distribution theory. *Papers in Regional Science* 5(1): 51–56.

Simini, F.; González, M. C.; Maritan, A.; and Barabási, A.-L. 2012. A universal model for mobility and migration patterns. *Nature* 484(7392): 96.

Spadon, G.; de Carvalho, A. C.; Rodrigues-Jr, J. F.; and Alves, L. G. 2019. Reconstructing commuters network using machine learning and urban indicators. *Scientific reports* 9(1): 1–13.

Stouffer, S. A. 1940. Intervening opportunities: a theory relating mobility and distance. *American sociological review* 5(6): 845–867.

Yang, Y.; Herrera, C.; Eagle, N.; and González, M. C. 2014. Limits of predictability in commuting flows in the absence of data for calibration. *Scientific reports* 4: 5662.

Zhang, N.; and Alipour, A. 2019. Integrated framework for risk and resilience assessment of the road network under inland flooding. *Transportation research record* 2673(12): 182–190.

Zipf, G. K. 1946. The P 1 P 2/D hypothesis: on the intercity movement of persons. *American sociological review* 11(6): 677–686.