

On Influencing Individual Behavior for Reducing Transportation Energy Expenditure in a Large Population

Shiwali Mohan¹, Frances Yan¹, Victoria Bellotti¹, Ahmed Elbery², Hesham Rakha² and Matthew Klenk¹

1. Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, California

2. Virginia Tech Transportation Institute, Transportation Research Plaza Blacksburg, Virginia

Abstract

Our research aims at developing intelligent systems to reduce the transportation-related energy expenditure of a large city by influencing individual behavior. We introduce COPTER - an intelligent travel assistant that evaluates multi-modal travel alternatives to find a plan that is *acceptable* to a person given their context and preferences. We propose a formulation for *acceptable* planning that brings together ideas from AI, machine learning, and economics. This formulation has been incorporated in COPTER producing acceptable plans in real-time. We adopt a novel empirical evaluation framework that combines human decision data with high-fidelity simulation to demonstrate a 4% energy reduction and 20% delay reduction in a realistic deployment scenario in Los Angeles, California, USA.

Introduction

Transportation is one of the largest consumers of energy in the world - in the United States, it accounted for 29% of energy consumption in 2016¹. Despite occupying a central place in the economy, transportation is far from efficient. Many areas of urban transportation networks are underutilized while other areas are congested. Congestion alone in the United States wastes 6.9 billion hours and 3.1 billion gallons of fuel per year (Schrang et al. 2015). Thus, transportation has become an important policy & innovation problem. Recently, public and private entities have begun introducing new transportation services including bike/scooter share, car share, ride hailing, and carpooling to complement private vehicles and existing public transit. This increasing accessibility of novel modes and the corresponding need for planning solutions creates an opportunity for *mobility marketplaces* that aggregate a variety of offerings into a single market, similar to what Travelocity or Expedia does for air travel. Such mobility marketplaces have the potential to influence people to adopt sustainable modes at a large scale.

Our long-term research goal is to develop AI methods to support efficient transportation reducing overall energy consumption and increasing citizen satisfaction in metropolitan areas. There are two important components to this problem of changing the behavior of a large population toward this *social good* goal. The first - the *multi-modal route planning* problem - addresses how to reason about the large decision

space of a variety of modes that are accessible to a person to determine an energy-efficient trip. The second component - the *influence* problem - is guiding an individual's behavior to actually adopt trips that reduce the overall energy expenditure of a region. While the multi-modal route planning problem has been studied in detail by the AI community (Bast et al. 2016), the influence problem is largely unexplored.

Our approach to the influence problem is embedded within an intelligent *travel assistant* - Collaborative Optimization and Planning for Transportation Energy Reduction (COPTER). We illustrate COPTER's desired behavior with the following example. Assume that John makes a regular trip driving to his office every weekday morning and has COPTER installed on his smartphone. In a *timely* fashion, 15 minutes before he has to leave, COPTER suggests an alternative trip - walk to the bus stop and take the direct bus to office - chosen from all available alternatives. This trip is not only energy-efficient, but also *acceptable* to John given the context of his travel. COPTER makes this suggestion by considering that the route is direct, that a lot of people in his neighborhood take the bus, and that he is traveling to his workplace and therefore, will not be carrying heavy loads. Further, COPTER recognizes that he cares about his impact on the environment and frames this information in a *compelling* way by telling him how much emissions he can reduce by taking this trip. There are several aspects to influencing someone to change their behavior. In this example, reasoning about acceptability of a trip recommendation, when should it be communicated, and how it should be expressed is useful for influencing John.

In this paper, we study a specific aspect of the influence problem - how a person's context and preferences can be integrated into the planning framework to produce an *acceptable* recommendation and how this acceptability impacts a person's adoption of the recommendation. To do this, we bring together ideas from AI planning, machine learning, as well as economic choice theory. More specifically, our paper makes the following contributions:

1. Introduces the transportation influence problem,
2. Proposes and implements a novel integration of AI planning theory and economic choice theory,
3. Defines a machine learning model to capture the acceptability of recommended plans, and
4. Evaluates the potential impact on Los Angeles by combining human decision data with a high-fidelity transportation simulation model.

¹www.eia.gov

The Setup

To understand the novelty of our approach, we first review the two strands of research we build upon, and later, we describe their limitations for the influence problem.

Multi-Modal Route Planning

Routing in transportation networks is a well studied problem dating back to Dijkstra's algorithm (Dijkstra 1959). Today, through pre-computation and heuristics, planners identify optimal routes for continent-size networks in milliseconds (Bast et al. 2016).

These approaches use the standard transportation network representation as a directed graph $G = (V, E)$, where V is a set of nodes and $E \subseteq V \times V$ is a set of directed edges. The formalism also includes a standard representation of time, T . A request for transporting a person from their origin location $v \in V$ starting at a certain time $t_s \in T$ to their destination $w \in V$ before a certain deadline $t_e \in T$ is formalized as $r = (v, w, t_s, t_e)$. The planning problem is formalized as (G, r) . Traversal of each edge is represented as a pair (e, t) where a person begin moving along the edge e at time t . A plan $\pi = ((e_1, t_1), \dots, (e_n, t_n))$ is a valid solution of the problem (G, r) if and only if:

1. the edges connect the origin with the destination. $e_1 = (v_0, v_1), \dots, e_n = (v_{n-1}, v_n) \wedge (v_0 = v) \wedge (v_n = w)$,
2. temporal constraints are satisfied, $t_s \leq t_1 + dur(e_1) \leq t_2 + dur(e_2) \leq \dots \leq t_n + dur(e_n) \leq t_e$ where $dur(e)$ represents the time it takes to traverse the edge e .

To find an optimal plan, cost function $cost(e)$ is defined for every edge in G . An optimal plan is a plan that minimizes the cumulative cost $\pi^* = \arg \min_{\pi \in \Pi} \sum_{e \in \pi} cost(e_i)$. In the simple case, the cost is the time it takes to traverse the edge, $cost(e_i) = dur(e_i)$. A* search with a reasonable heuristic (e.g., euclidean distance) efficiently solves these problems.

Recent work provides extensions for multi-modal trips (Dibbelt, Pajor, and Wagner 2015) and more elaborate cost functions. This is achieved by adding additional edges between vertices corresponding to the various modes (e.g., public transportation, biking, walking) that are available to traverse them. When the edges represent public transit vehicles, they have temporal constraints representing the vehicle schedules. Each edge e is labeled with the mode it represents, e.g., $lbl(e) = b$ for an edge representing taking the bus. Consequently, each plan π corresponds to a sequence of labels, $word(w_p)$. This supports an additional requirement on valid solutions that the word is part of the language of acceptable mode sequences ($w_p \in L$), represented using regular expressions. This constraint eliminates plans that are impossible, such as taking a bus and then getting on a bike. Others (Botea 2016) have shown that in this formulation, the cost function can consider additional dimensions including price, energy, and schedule risk. The cost function is formalized as $cost(\pi) = \sum_{e \in \pi} \sum_{\varphi \in \Phi} \theta_\varphi \varphi(e)$ where Φ is the set of evaluative functions including energy consumed, price, etc. and θ_φ corresponds to how much importance that evaluative function has toward computing the overall edge cost.

Utility-Based, Probabilistic Choice Theory

Economists use the framework of rational choice theory (Tversky and Kahneman 1986) to study how humans make

choices from a set of competing alternatives. Empirical research in transportation economics (Domencich and McFadden 1975) has shown that such frameworks are useful in predicting transportation mode shares in urban populations. The theory posits a choice process: for a trip, the individual first determines the available mode alternatives; next, evaluates the attributes of each alternative relevant to the travel decision, and then, uses a decision rule to select an alternative. Rational decision making perspective motivates *utility maximization* as a relevant decision rule. Utility maximization is based on two fundamental concepts. First, the attributes of a mode alternative can be reduced to a scalar utility value. Second, the decision maker selects the alternative with the highest utility value. The deterministic theory assumes all aspects of the decision process can be observed and measured. Transportation decision-making is only partially-observable. By adding noise, the utility of an alternative for an individual can be split into two components - observable utility and error:

$$util(x_i, f_p) = val(x_i, f_p) + \epsilon(x_i, f_p)$$

Here $val(x_i, f_p)$ is the observable portion of the utility. This quantity is modeled as a linear combination of measurable attributes such as travel time, travel cost, walk distance, etc:

$$val(x_i, p) = \gamma_1 \times x_{i1} + \gamma_2 \times x_{i2} + \dots + \gamma_k \times x_{ik} + \lambda_1 \times f_{p1} + \lambda_2 \times f_{p2} + \dots + \lambda_l \times f_{pl} \quad (1)$$

γ_k is the parameter which defines the direction and importance of the attribute k on the utility of an alternative i . Similarly, λ_l defines the direction and importance of at the person's characteristic l .

$\epsilon(x_i, f_p)$ is the unobservable portion - the error term - of the utility function. This theory assumes that the error terms of various alternatives are independent. These assumptions about the error term lead to a *multinomial logit model* that estimates the probability $Pr(i, p)$ of the person p selecting the alternative i from the set of alternatives C :

$$Pr(i, p) = \frac{e^{val(x_i, f_p)}}{\sum_{j \in C} e^{val(x_j, f_p)}} \quad (2)$$

Empirical work in transportation research (Ben-Akiva, Lerman, and Lerman 1985) has studied how γ and λ parameters in equation 1 can be estimated from human choice data using maximum likelihood estimation given equation 2.

Acceptable Planning

There are two main challenges that have to be addressed in order to determine trip recommendations that are *acceptable* - useful in influencing behavior. The first one is that none of the theories described in the previous section explicitly provide a definition of *acceptability*. Planning computes the optimal path for an individual given a cost function defined for every edge while the choice theory predicts the probability of a trip occurring given some observable characteristics about the person and the mode of transport. In COPTER, each traveler has an expected plan for an upcoming trip, and our goal is to produce an *acceptable* alternative that changes their behavior to reduce energy consumption across the network. The second is an integration issue. For efficient search, the planning theory requires the costs to be

defined locally at each edge. The choice theory and empirical evidence from transportation research on the other hand suggests that there are aspects of a person’s trip context that influence mode choice but cannot be defined over edges in a graph. These include the household income, number of vehicles at home, etc. In the sections below, we propose how these challenges can be addressed for *acceptable* planning.

Defining Acceptability

We adapt choice theory’s decision process to incorporate *acceptability* as follows. John (from the example in the Introduction) drives to his office every weekday on a certain route. The choice theory suggests that this usual route has a measurable utility, $val(x_u, f_p)$ and a probability of being selected based on equation 2. Here x_u is the vector of attributes representing John’s usual mode such as time, distance, etc. and f_p is a vector that describes John’s attributes that pertain to mode selection such as income, education level etc. Upon receiving the recommendation, r , from COPTER, John evaluates it against his usual means of travel. The recommendation r has a measurable utility, $val(x_r, f_p)$, given the utility function underlying John’s preferences. On adopting the recommendation, John will experience a *switching cost* - $\Delta_{u,r}$ - the difference in utilities of his usual route and the proposed route. Conversely, $\Delta_{r,u} = -\Delta_{u,r}$, can be understood as the *switching gain* a person makes on adopting the recommendation r .

Following probabilistic choice theory (equation 2):

$$\frac{e^{val(x_u, f_p)}}{e^{val(x_r, f_p)}} = \frac{Pr(u, p)}{Pr(r, p)}$$

$$\Delta_{u,r} = val(x_u, f_p) - val(x_r, f_p) = \ln \frac{Pr(u, p)}{Pr(r, p)} \quad (3)$$

$$\Delta_{r,u} = -\Delta_{u,r} = \ln \frac{Pr(r, p)}{Pr(u, p)}$$

We postulate that the acceptability of a route is related to its switching gain. If we can estimate the likelihood of a person taking their usual mode of travel and that of the recommended mode, we can compute the change in utility and consequently estimate the recommendation’s acceptability. This formulation is a significant departure from prior empirical work in transportation research. The prior work focuses on estimating the coefficients in equation 1 to compute exact utility values given measurements of relevant, observable attributes. For planning, this exact computation of utility value is not necessary and the switching gain (and acceptability) can be computed from estimated likelihoods alone.

Estimating Switching Gain

As explained earlier, in order to estimate the acceptability of a recommendation r for a person p , it is crucial to estimate the likelihood $Pr(r, p)$ of it being selected as well as for that person’s usual mode of travel, $Pr(u, p)$. We adopt a data-driven, ML-based approach to estimate these likelihoods. A typical ML classifier generates a prediction of a class, given a set of input features, by computing likelihoods for each class and selecting the class with the highest likelihood.

To develop our ML classifiers, we used the California Household Travel Survey (CHTS) that consists of a single-day (distributed in a year) travel diary of people from 58

Mode	Baseline 1	Baseline 2	RF	MLP
<i>Walk</i>	0.00	0.12	0.82*	0.62
<i>Cycle</i>	0.00	0.00	0.81*	0.28
<i>Bus</i>	0.00	0.02	0.78*	0.38
<i>Subway/train</i>	0.00	0.00	0.58*	0.05
<i>Drive</i>	0.72	0.56	0.93*	0.86
<i>Ride</i>	0.00	0.28	0.84*	0.65
<i>Motorcycle</i>	0.00	0.00	0.80*	0.00
Total	0.68	0.40	0.88*	0.74

Category	Baseline 1	Baseline 2	RF	MLP
<i>Non-motorized</i>	0.00	0.05	0.83*	0.60
<i>Public transit</i>	0.00	0.14	0.79*	0.43
<i>Motorized</i>	0.90	0.82	0.97*	0.93
Total	0.68	0.70	0.94*	0.86

Table 1: F1-scores for predicting mode classes and categories with various ML methods. Baseline 1 is most frequent, Baseline 2 is weighted random, RF is random forest classifier, and MLP is multi-layer perceptron

counties of California and 3 counties of Nevada. To ground our research within the context of a specific metropolitan area, we extracted trip data for 4,889 Los Angeles (LA) residents (from 2,006 households). Our trip dataset contains 78,380 trips undertaken in LA county. For each trip, the data consists of the main mode of transport as well as other demographic and context information.

Features and Classes Our trips dataset includes trips that were undertaken using 7 different modes including - walk, cycle, bus, subway/train, drive, ride, motorcycle. For each trip, we also extracted these features:

- *trip-related*: trip distance,
- *demographics*: such as education level, number of people/students/workers in the household,
- *employment*: number of hours worked every week, income, number of jobs, flexibility of work,
- *mode accessibility*: number of automobiles in the household, number of bicycles in the household, owns a driver’s license & transit pass, and
- *mode experience*: transit used, bike trips made & walking trips made in the past week.

Note that these features correspond to x and f_p in the choice theory formulation.

ML Models & Results We explored two different multi-class prediction problems: predicting the mode and predicting the category of the mode where we categorized walking & cycle as *non-motorized*, driving, riding &, motorcycle as *motorized*, and bus & subway as *public transit*. For each problem, we use two different algorithms: a random forest classifier with 20 estimators of 30 depth each; and a multi-layer perceptron with 4 layers ([1000, 500, 100, 100]). We found these parameters through a bucketed linear hyperparameter search. Classifier performance was compared with two baselines: assigning the most frequent class to every sample (*most frequent*); and distributing the labels by their frequency in the dataset (*weighted random*).

F1 scores for both classifiers as well as baselines on a 20% test set are shown in Table 1. We see that both ML

classifiers perform significantly better than the baselines for not only predicting each class (& category) but have a better overall performance as well. This suggests that the ML classifiers contain useful predictive information that can be used to estimate the switching gains as well as acceptability of a recommendation. We further see that the random forest classifier has better performance than the multi-layer perceptron. This is expected due to the limited size of the dataset. We further see that the performance on category prediction is better than mode prediction. Again, this is to be expected as combining mode classes leads to more data to learn from and consequently better performance.

Switching Gain By using the ML model described here to estimate likelihoods of various modes, we can compute the switching gain to any mode from a specific mode a person usually takes. The ML models use a diverse types of information including demographics, trip features, and mode experience with modes to estimate likelihoods and consequently, switching gain is personalized to an individual's context. Intuitively, switching gain captures the idea that it is easier to switch to modes that are used by people who are similar to you and make similar trips.

Evaluating Acceptability

To assess the impact on switching gain (acceptability) on likelihood of recommendation adoption, we conducted a mode adoption study with 49 (27 female, 22 male) participants who were regular drivers in Los Angeles. The participants were gathered by a paid recruiter who ensured compliance. The participants were provided \$30 for completing the study.

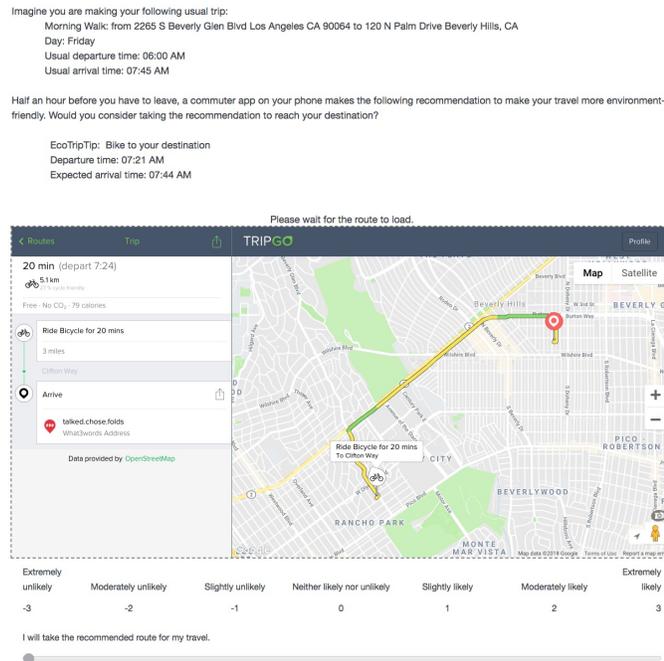


Figure 1: An example of the mode adoption question in the survey along with the response measurement instrument.

Materials Each study participant took two surveys:

1. A *profiler* survey included questions for features used in our classifier. For each participant, the survey collected a set of trips they regularly made by driving in terms of the origin, destination, departure time, and purpose.
2. A *mode adoption* survey presented each participant 10 influence questions. Each question proposed an alternative mode route for one of their typical driving trips. Care was taken to not include implausible alternatives such as walking a 10 mile distance. Each question included information about mode characteristics such as the time of departure and a visual representation of the proposed route (shown in Figure 1). The participants indicated likelihood of adopting the recommendation using a 7-point Likert scale (from *extremely unlikely* (1) to *extremely likely* (7)).

Stated preference methods are typically used to measure the acceptability of new products and provide realistic estimates of population behavior (Johnston et al. 2017).

Models The study measures mode adoption as an ordinal variable on the Likert scale. To simulate real-world behavior where someone will either follow the recommendation or not, we compressed the ordinal response by categorizing everything reported as *slightly likely* or higher as 1 and others at 0. Our assessment of the impact of switching gain on adoption included the following regression models:

1. For the ordinal variable, we fit a mixed-effects linear regression of the form $y = \alpha + \beta x + \gamma z + \epsilon$, where y is the adoption vector, x is a fixed-effect vector corresponding to an independent variable of interest, and z is a random effect vector corresponding to a participant. α is the intercept and β and γ are coefficients and ϵ is the error term.
2. For the binary adoption variable, we fit a mixed-effects logistic regression of the form $Pr(y) = 1/(1 + e^{-(\alpha + \beta x + \gamma z + \epsilon)})$, where $Pr(y)$ is the probability of observing the specific adoption value (0/1), other coefficients are similar to those in the previous model.

Participants were included as random effects in these model to account for individual differences in interpretation of the scale. For independent variables (x , in the models above), we explored three definitions of acceptability:

1. switching gain, $\Delta_{r,u}$
2. odds, $e^{\Delta_{r,u}} = Pr(r, p) / Pr(u, p)$
3. probability, $Pr(r, p)$

where r is the recommended mode, u is the usual mode for person p and $Pr()$ is the likelihood from our ML models.

Results Overall, the results show that people reported a willingness to change their travel behavior. Detailed results are shown in Table 2. For each independent variable, we report R^2c , the proportion of the data that is explained by the complete model, and R^2m , the proportion is explained by the variable. We see that all definitions of acceptability significantly impact the adoption of recommendation. The coefficients align with the intuition about the domain - as the cost of adopting the recommendation decreases (switching gain increases), adoption increases. We see that R^2c , R^2m are higher for odds. These results suggest that it is the best predictor mode adoption because it explains the most data. It is noteworthy that odds and probability have similar impact on adoption. This is likely due to our participants being regular drivers with $Pr(u, p) \approx 1$. With a more diverse population in a larger deployment, the differences between the

definitions can be teased apart.

Dependent variables → Independent variables ↓	Adoption (ordinal)	Adoption (binary)
(intercept)	-0.017	-0.185
switching gain, $\Delta_{r,u}$	0.108*	0.104*
R^2_m	0.034	0.035
R^2_c	0.347	0.270
(intercept)	-0.949	-1.065
odds, $e^{\Delta_{r,u}}$	2.386***	2.159*
R^2_m	0.075	0.064
R^2_c	0.379	0.300
(intercept)	-0.964	-1.080
probability, $Pr(r,p)$	3.623***	3.317*
R^2_m	0.066	0.058
R^2_c	0.369	0.293

Table 2: Regression modeling results for ordinal and binary mode adoption. *** $p < 0.001$, ** $p < 0.05$, * $p < 0.1$

Estimating Likelihood of Adoption

To identify the recommendation that reduces the most expected energy, it is necessary to estimate the likelihood of adoption. Recall that our motivating example suggests that a recommendation should be *timely*, *acceptable*, and *compelling* to motivate behavior change. At this point, we have only evaluated acceptability. Therefore, COPTER uses the logit model of the odds to estimate the probability of a person adopting a recommended route.

Finding Energy-Efficient Acceptable Plans

Figure 2 provides an overview of COPTER’s operation. In a deployment, we expect to maintain a series of future trips for each user. Before an expected departure, COPTER computes an energy efficient *acceptable* alternative and sends it to the user 15 minutes before they would have to leave for the new trip. To compute this plan, we augment the multi-modal planning formulation with our acceptability model:

1. Generate a mode candidate set M_p for the person p for whom the request is made for. M_p can be generated by acquiring information about if they can walk and if they have a bike. For example, someone who doesn’t own a bike but can walk $M = \{walk(w), bus(b), subway(s)\}$.
2. Given M_p , determine the language L_p . Recall that a language is a set of regular expressions that denote valid multi-modal plans for this traveler. For the example above, $L_p = \{w*, w * b + w*, w * s + w*\}$, where w is *walking*, b is *bus*, and s is *subway* or *train*. Note that while this is a fairly simplistic set, the formulation can be extended to include more complex plans.
3. For every element in L_p , generate the most time-efficient plan (using the multi-modal formulation) where $cost(e) = dur(e, m)$ where m is the mode used to traverse the edge. The heuristic is the euclidean distance. This process will generate a candidate set of plans, Π_p .
4. Compute the energy reduction in each plan $\pi \in \Pi_p$ using existing mesoscopic energy models (Elbery et al. 2018).
5. Evaluate the likelihood of adoption, $adopt(\pi, p)$, for every $\pi \in \Pi$ using the logit model of adoption.
6. Select a plan that has maximal expected energy savings, $\pi^* = \arg \max_{\pi \in \Pi} adopt(\pi, p) \times energy(\pi)$.

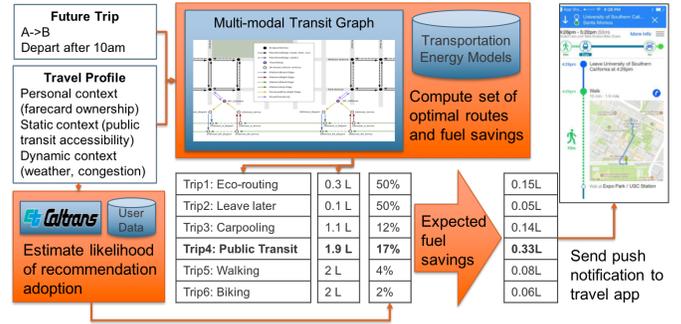


Figure 2: COPTER identifies the recommendation with the largest expected reduction in energy consumption by acceptable planning and energy modeling.

Simulation Experiments

While deploying COPTER in a large city to measure the impact is the ideal evaluation, it is extremely resource intensive to conduct. Without being incorporated into a widely used product, it would not be possible to iteratively evaluate different components. Therefore, we evaluate the potential impacts of COPTER using a high-fidelity simulation LA. In our deployment scenario, we assume that 10% of peak period drivers use an app (e.g., Waze, TripGo) through which we can offer alternative suggestions.² Our study area includes diverse offerings with over 170,000 roadway links and over 1 million daily transit trips. During the AM period (7am-10am) and PM period (4pm-7pm), this area includes 1.2 million and 1.7 vehicle trips respectively. As the effects of congestion continue past the peak period, our study considers an AM period (7am-12pm) and a PM period (4pm-9pm).

Simulation Model

To measure the potential impact of influencing behavior, we combined the results of the human study with a state-of-the-art simulation model of the LA region (Elbery et al. 2018). The simulation model supports city-level networks with different modes of transportation including cars, buses, railways, walking, biking, and carpooling. It utilizes both microscopic and mesoscopic simulation to leverage their respective strengths of accuracy and scalability. The simulation spatially partitions the road network enabling distinct portions of the region to be simulated in parallel. Additional parallel simulations capture loosely interacting modes (e.g., the rail and road network). This enables micro-simulation (i.e., considering driver behavior w.r.t. throttle position, braking, and lane choice every 0.1s) of the arterial and highway roads encoded as 10,650 links. Dynamic calibration aligns previous government planning documents with observed vehicle flows to capture the changing rates of vehicles entering the different networks throughout the day (Du et al. 2018). The model includes validated driver behavior (Rakha, Pasumarthy, and Adjerid 2009) and fuel models (Fiori, Ahn, and Rakha 2016). Therefore, the model provides accurate individual and aggregate predictions in terms of travel time and energy consumption.

Given our deployment assumption, we sample 10% out of the peak traffic periods to be influenced (7am-10am and

²This percentage was selected through negotiations with business groups and public sector stakeholders

	Baseline	Influence	Change (CI)
Total	3,195,637	3,048,278	-4.6%
Fuel (l)			(-3.6% -5.6%)
Total	249,221	199,395	-20%
Delay (hr)			(-13.6% -26.4%)

	Baseline	Influence	Change (CI)
Total	3,487,982	3,367,675	-3.5%
Fuel (l)			(-2.6% -4.3%)
Total	375,137	322,228	-14.1
Delay (hr)			(-10% -18%)

Table 3: Results of departures: (left) 7:00 - 12:00 (right) 16:00 - 21:00. The mean for each condition is reported along with the percent change in the Influence condition. We also report a 95% confidence interval for the difference in mean.

4pm-7pm) for our influenced population. Our experiment has two conditions:

1. Baseline: Influenced traveler drives their original route
2. Influence: Using the likelihood of adoption from the previous section, the influenced traveler either takes the maximum expected energy reduction route as determined by COPTER or drives their original route.

To account for the uncertainty in alternative adoption and background traffic rates, we ran each condition four to seven times. The simulation model runs in just over real time and therefore, a running a larger number of trials is prohibitive. 5 – 10 trials is standard practice for similar transportation studies. We report 95% confidence intervals for difference of means between the Baseline and Influence trials for total fuel consumption and congestion related delay. We also report the resulting changes in mode of the influenced travelers.

Results Table 3 contain the results for AM and PM periods. These results indicate a potential 4.6% energy reduction during the AM period and 3.5% energy reduction in the PM period. Congestion-induced delay results were even greater at 20% and 14%, respectively. This is significant considering the influenced population represents about 6% of the total traffic over the period. To identify an upper bound, we ran simulations in which every influenced traveler walked. In these simulations, we observed a reduction in energy consumption of 9% and a reduction in delay of 30% across the AM and PM periods. Table 4 indicates that almost half the influenced population is willing to take alternative modes. These mode switches are based on mode adoption study and maybe optimistic. Future work will refine these estimates through a deployment and measuring people’s behavior.

Mode	AM Share	PM Share
Car	54%	53%
Walk	42.7%	42.8%
Bike	3.6%	3.8%
Bus	38.9%	39.1%
Train	14.4%	14%

Table 4: Share of influenced population that used mode on trip. Sum is over 100% due to trips using multiple modes.

Related Work

The majority of transportation policy either focuses on systemic changes that are applied evenly to the entire population such as congestion pricing/dynamic tolling (Brownstone et al. 2003) or studies long term behavior change in individuals to encourage sustainable commuting (Castellani et al. 2016). Recently, there has been a shift toward personalizing the interventions to increase the use of sustainable modes. IncenTrip (Sun and Zhang 2018) allocates incentives from government sponsored programs such that each dollar spent provides the largest environmental impact. TriPod

(Azevedo et al. 2018) takes a different approach in which companies give travelers points for alternative transportation choices to be redeemed at local businesses. While these systems primarily rely on monetary incentives to affect travel behavior, Metropia (Zhu, Hu, and Chiu 2018) has shown that congestion information alone can influence traveler departure time. Our work differs from the above by identifying a personalized recommendation from alternatives based on network conditions without any monetary incentives.

In addition to algorithmic work in route planning, the AI community also works on transportation prediction problems. Using position traces and other data, systems predict the mode and destination of travelers (Liao et al. 2007; Song, Kanasugi, and Shibasaki 2016). These works address prediction problems in isolation, and, as noted by AI and Society researchers (Inclezan and Prados 2017), it is how AI approaches integrated in smart cities applications that determine if social positive outcomes are realized. By working across disciplines, we study how energy consumption can be reduced through voluntary behavior changes. With respect to the transportation influence problem, COPTER could incorporate these isolated approaches to predicting future trips and monitor if travelers follows the recommendation.

Discussion and Conclusion

This paper presents COPTER, an intelligent travel assistant that is designed to influence individual travel behavior to reduce the total energy consumption of a large city. To do this, we propose an integrative framework that is a novel combination of multi-modal trip planning, transportation choice modeling, and machine learning. We introduce the notion of acceptability of a recommended alternative and define it in terms of the change in utility. We show how this change can be estimated using ML models trained on the CHTS dataset. We identify a useful definition of acceptability through a human mode choice study. Finally, in a simulation experiment over the Los Angeles area we show that if 10% peak travelers were receiving influence messages, estimated energy consumption would be reduced by 4.6% in the AM period and 3.5% in the PM period with corresponding reductions in congestion induced delay of 20% and 14%, respectively.

Future work will consider additional alternatives. Eco-routing feedback control selects driving routes based on real-time link estimates of fuel consumption (Elbery et al. 2018). Additionally, we would also like to explore techniques for departure time optimization as well as including carpooling in the set of alternatives. From a social good perspective, COPTER is ready for deployment with influenced messages sent through the TripGo application. The next step is to evaluate the influence model by measuring actual traveler behavior and then feeding those results back into our simulation study. Later work will define and assess the impacts of timeliness of recommendation and compelling messaging. Urban transportation is a huge problem for the environment

and citizen's quality of life. COPTER demonstrates how AI techniques can be adapted and combined with research from other disciplines to alleviate congestion and improve the environment through transportation influence.

Acknowledgements

This work was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0000612. The authors would like to thank the City of Los Angeles and CalTrans for providing data necessary for the modeling efforts.

References

- Azevedo, C. L.; Seshadri, R.; Gao, S.; Atasoy, B.; Akkinapally, A. P.; Christofa, E.; Zhao, F.; Trancik, J.; and Ben-Akiva, M. 2018. Tripod: Sustainable Travel Incentives with Prediction, Optimization, and Personalization. In *the 97th Annual Meeting of Transportation Research Board*.
- Bast, H.; Delling, D.; Goldberg, A.; Müller-Hannemann, M.; Pajor, T.; Sanders, P.; Wagner, D.; and Werneck, R. F. 2016. Route Planning in Transportation Networks. In *Algorithm Engineering*. Springer International Publishing. 19–80.
- Ben-Akiva, M. E.; Lerman, S. R.; and Lerman, S. R. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*, volume 9. MIT press.
- Botea, A. 2016. Hedging the Risk of Delays in Multimodal Journey Planning. *AI Magazine* 37(4).
- Brownstone, D.; Ghosh, A.; Golob, T. F.; Kazimi, C.; and Van Amelsfort, D. 2003. Drivers Willingness-to-pay to Reduce Travel Time: Evidence from the San Diego I-15 Congestion Pricing Project. *Transportation Research Part A: Policy and Practice* 37(4):373–387.
- Castellani, S.; Colombino, T.; Grasso, A.; and Mazzega, M. 2016. Understanding Commuting to Accompany Work Organisations' and Employees' Behaviour Change. In *Smart Cities Conference (ISC2), 2016 IEEE International*, 1–6. IEEE.
- Dibbelt, J.; Pajor, T.; and Wagner, D. 2015. User-Constrained Multimodal Route Planning. *Journal of Experimental Algorithmics* 19:3.2:1–3.2:19.
- Dijkstra, E. W. 1959. A Note on Two Problems in Connection with Graphs. *Numerische Mathematik* 1(1):269–271.
- Domencich, T. A., and McFadden, D. 1975. Urban Travel Demand-A Behavioral Analysis. Technical report, Transport and Road Research Laboratory.
- Du, J.; Elbery, A.; Rakha, H.; and Klenk, M. 2018. Microscopic Simulation and Calibration of a Large-Scale Metropolitan Network: Issues and Proposed Solutions. In *the 97th Annual Meeting of Transportation Research Board*.
- Elbery, A.; Dvorak, F.; Du, J.; Rakha, H.; and Klenk, M. 2018. Large-scale agent-based multi-modal modeling of transportation networks - system model and preliminary results. In *Conference: 4th International Conference on Vehicle Technology and Intelligent Transport Systems*, 103–112.
- Fiori, C.; Ahn, K.; and Rakha, H. A. 2016. Power-based electric vehicle energy consumption model: Model development and validation. *Applied Energy* 168:257–268.
- Incelezan, D., and Pradanos, L. I. 2017. A Critical View on Smart Cities and AI. *Journal of Artificial Intelligence Research* 60:681–686.
- Johnston, R. J.; Boyle, K. J.; Adamowicz, W.; Bennett, J.; Brouwer, R.; Cameron, T. A.; Hanemann, W. M.; Hanley, N.; Ryan, M.; Scarpa, R.; et al. 2017. Contemporary Guidance for Stated Preference Studies. *Journal of the Association of Environmental and Resource Economists* 4(2):319–405.
- Liao, L.; Patterson, D. J.; Fox, D.; and Kautz, H. 2007. Learning and Inferring Transportation Routines. *Artificial Intelligence* 171(5):311 – 331.
- Rakha, H.; Pasumarthy, P.; and Adjerid, S. 2009. A Simplified Behavioral Vehicle Longitudinal Motion Model. *Transportation letters* 1(2):95–110.
- Schrank, D.; Eisele, B.; Lomax, T.; and Bak, J. 2015. 2015 Urban Mobility Scorecard. *Annual Urban Mobility Scorecard*.
- Song, X.; Kanasugi, H.; and Shibasaki, R. 2016. Deep-Transport: Prediction and Simulation of Human Mobility and Transportation Mode at a Citywide Level. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 16, 2618–2624.
- Sun, Y., and Zhang, L. 2018. Microeconomic Model for Designing Public Transit Incentive Programs. *Transportation Research Record*.
- Tversky, A., and Kahneman, D. 1986. Rational Choice and the Framing of Decisions. *Journal of Business* S251–S278.
- Zhu, X.; Hu, X.; and Chiu, Y.-C. 2018. Will Information and Incentive Affect Travelers Day-to-Day Departure Time Decisions? An Empirical Study of the Decision Making Evolution Process. In *the 97th Annual Meeting of Transportation Research Board*.