

# Acceptable Planning: Influencing Individual Behavior to Reduce Transportation Energy Expenditure of a City

**Shiwali Mohan**

SHIWALI.MOHAN@PARC.COM

*Palo Alto Research Center, Mail Stop: 3333  
Coyote Hill Road, Palo Alto, CA 94034 USA*

**Hesham Rakha**

HRAKHA@VTI.VT.EDU

*Virginia Tech Transportation Institute, Mail Stop: 3500  
Transportation Research Plaza Blacksburg, VA 24061*

**Matthew Klenk**

MATTHEW.KLENK@PARC.COM

*Palo Alto Research Center, Mail Stop: 3333  
Coyote Hill Road, Palo Alto, CA 94034 USA*

## 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 that produces acceptable plans in real-time. We adopt a novel empirical evaluation framework that combines human decision data with a high fidelity multi-modal transportation simulation to demonstrate a 4% energy reduction and 20% delay reduction in a realistic deployment scenario in Los Angeles, California, USA.

## 1. 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 2016 (EIA, 2016). 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 (Schrank, Eisele, Lomax, & Bak, 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 enable the creation of new technologies to assist travelers in making decisions that meet their needs and benefit their community.

Our long-term research goal is to develop AI methods to support efficient transportation decision-making reducing overall energy consumption and increasing citizen satisfaction in metropolitan areas. There are two important components to changing the transportation 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, routes, and combinations that are accessible to a person to determine an energy-efficient trip. The second component - the *influence* problem - is guiding an individual's behavior such that they 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 in the AI literature. When influence is considered, it is from the perspective of monetary incentives (Azevedo et al., 2018; Sun & Zhang, 2018), which poses a business challenge. Here we focus on understanding the factors that underlie transportation decision-making to enable influence without monetary incentives.

Our research on AI systems that act to influence human behavior towards a social good goal is part of a larger research agenda on design and analysis of *human-aware* AI systems. Often while designing AI systems, little attention is devoted to modeling the behavior of humans who invariably are crucial decision makers. Human-aware AI systems (Kambhampati, 2019) seek to address this gap and pose modeling humans as a question central to AI system design. While this is a worthwhile pursuit, current human modeling methods in AI are extremely limited in their scope (Albrecht & Stone, 2018). Largely, human modeling research in AI originates in game domains where humans are modeled either as collaborators or opponents with fixed objective functions. These models are not sufficient for representing the causation of dynamic and evolving human behavior in the real-world (such as choosing a mode of transport) and consequently, cannot be applied to influence it.

Our approach to building human-aware AI systems is to borrow models and methods from social sciences, in particular, behavioral economics, and explore how they can be integrated in computational systems in a principled fashion. Additionally, we relinquish typical AI metrics of accuracy and efficiency and define new *human-centered* desiderata for AI systems to ensure that we are making progress towards systems that are truly human-aware. In our problem, the desiderata is related to *influence* - how likely is a person to change their behavior when provided with a recommendation. To develop a computational formulation of *influence*, we engage with research in economics on rational choice and human preference modeling. Further, we adopt empirical methods prevalent in social science research to evaluate our proposed formulation. Then, we engage with transportation modeling research to simulate the impact of applying our approach on the energy consumption of a large city. By engaging these disciplines, we propose a comprehensive methodology for developing intelligent solutions for social good. More specifically, our paper makes the following contributions:

1. Introduces the transportation influence problem for social good and a useful component - *acceptability*,
2. Develops a machine learning approach to estimate acceptability,
3. Evaluates the formulation of acceptability using a human participant study,

4. Proposes and implements a novel integration of AI planning theory and economic choice theory for *acceptable planning*, and
5. Evaluates the integrated approach by combining human decision data with a high fidelity transportation simulation model of the Los Angeles area.

## 2. The Traveller Influence Problem

Social sciences research proposes several factors that are useful for influencing someone to change their behavior. The transportation mode choice framework (Ben-Akiva, Lerman, & Lerman, 1985) suggests that attributes that define the context in which the trip is undertaken, such as distance of the trip, cost of taking the mode, duration of travel etc. play a role in people’s preference for a specific mode. Recommending a mode that is higher on a person’s preference list may increase the likelihood that the recommendation will be adopted. Just-in-time adaptive interventions for health behavior change (Nahum-Shani, Hekler, & Spruijt-Metz, 2015) suggest that recommendations provided at opportune times are more likely to influence behavior. Further, research in persuasive messaging (Benoit & Benoit, 2008) has studied the impact of various types of message framing - how information is conveyed - has on individuals in influencing their choices and behaviors. Building on these results, we define the travel influence problem as delivering an *acceptable, timely, and compelling* travel recommendation that results in a person selecting a sustainable mode.

Our approach to the influence problem is embedded within an intelligent *travel assistant* - Collaborative Optimization and Planning for Transportation Energy Reduction (COPTER) and aims at operationalizing influence strategies for travellers. Figure 1 contains a functional description of the system and we illustrate COPTER’s desired influence behavior with the following example. Assume that Jane makes a regular trip driving to her office every weekday morning and has COPTER installed on her smartphone. In a *timely* fashion, 15 minutes before she has to leave, COPTER suggests an alternative trip - walk to the bus stop and take the direct bus to her office - chosen from all alternatives available to Jane. This trip is not only energy-efficient, but also *acceptable* to Jane given the context of her travel and her preferences. COPTER makes this suggestion by considering that the route is direct, that a lot of people in her neighborhood take the bus, and that she is traveling to her workplace and therefore, will not be carrying heavy loads. Further, COPTER recognizes that she cares about her impact on the environment and frames this information in a *compelling* way by telling her the emissions she can reduce by taking this trip.

This paper explores how acceptability of a mode can be defined for transportation planning and how acceptability affects the adoption of a recommendation. The remainder of the paper is organized as follows. We begin by describing the multi-modal planning framework and choice theory framework that form the basis of our approach, in Section 3. Then, in Section 4, we formally define the notion of *acceptability* for planning. We develop a machine learning method to estimate individual acceptability from a population-level, transportation behavior dataset from Los Angeles, California in Section 5. We show that this acceptability model is predictive of adoption through a human participant study in Section 6. Upon validating our definition of acceptability, we incorporate it in the planning framework in Section 7. Finally, we demonstrate the energy impact of our approach using a high-fidelity simulation model of transportation energy integrated with human behavior

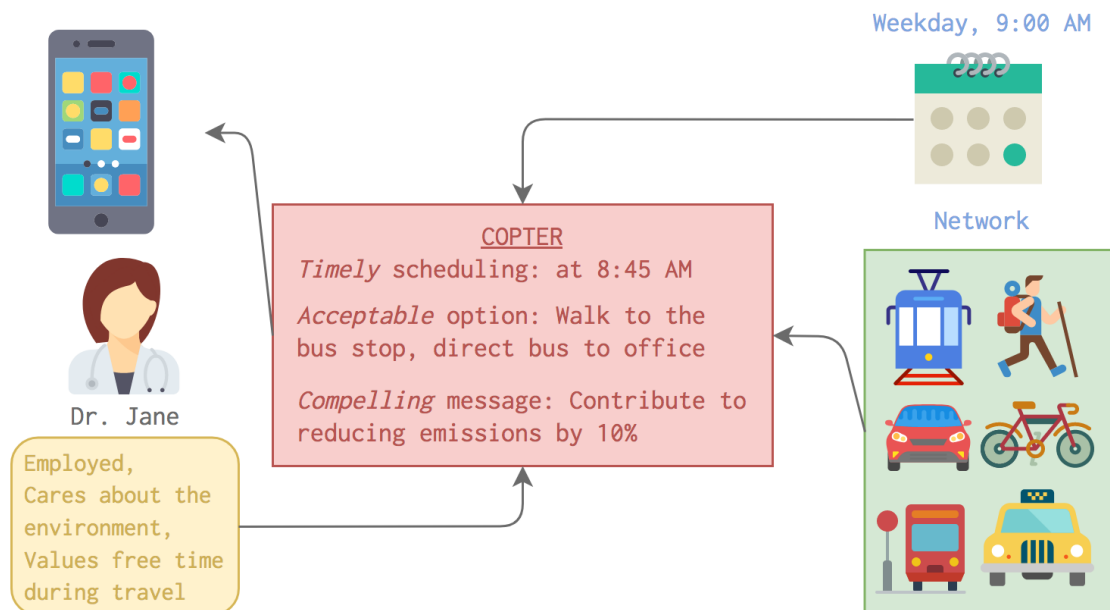


Figure 1: The transportation influence problem seeks to alter traveler behavior by suggesting routes based on the available options, the dynamic context including trip purpose, and a model of the traveler’s preferences.

in Section 8. In Section 9, we study research threads pursued by other researchers from AI, transportation, and behavior change. We conclude with a discussion of the future work in Section 10.

### 3. Preliminaries

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 transportation influence problem.

#### 3.1 Multi-Modal Route Planning

Routing in transportation networks is a well studied problem in computer science dating back to Dijkstra’s algorithm (Dijkstra, 1959). Today, through pre-computation and heuristics, planners identify optimal routes for continent-size networks while guaranteeing optimal solutions in fractions of seconds (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  $q = (v, w, t_s, t_e)$ . The planning problem is formalized as  $(G, q)$ . Traversal of each edge is represented as a pair  $(e, t)$  where

a person begins 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  $(G, q)$  if and only if:

1. the edges connect the origin with the destination.  $e_1 = (v_0, v_1), e_2 = (v_1, v_2), \dots, e_n = (v_{n-1}, v_n) \wedge (v_0 = v) \wedge (v_n = w)$ ,
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  $\pi^* = \arg \min_{\pi \in \Pi} \sum_{e \in \pi} cost(e_i)$ , where  $\Pi$  is the set of all valid plans. 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, & 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  $\pi$  corresponds to a sequence of labels,  $word(w_p)$ . This imposes a third 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; Dellling, Dibbelt, Pajor, Wagner, & Werneck, 2013; Dvorak, Mohan, Bellotti, & Klenk, 2018) 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  $\Phi$  is the set of evaluative functions including energy consumed, price, etc. and  $\theta_\varphi$  corresponds to how much importance that evaluative function has toward computing the overall edge cost.

To support personalization, numerous researchers have extended the basic notion of an edge-based objective function to capture more information. For example, deployed planners now combine travel time, cost, and convenience measures in selecting multi-modal plans (Dellling et al., 2013). In the research setting, methods for incorporating even more criteria (e.g., robustness to missing connections, Botea, 2016, safety, X. Liu, Fritz, and Klenk, 2018) have been used.

### 3.2 Utility-Based, Probabilistic Choice Theory

Economists use the framework of rational choice theory (Tversky & Kahneman, 1986) to study how humans make decisions from a set of competing alternatives. Empirical research in transportation economics (Domencich & 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. The 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. This assumption implies a compensatory decision process; i.e, the decision maker will make *trade-offs* among the various attributes. For example, an individual may choose a pricier mode if the comfort provided by that mode compensates for the increased cost. Second, the decision maker selects the alternative with the highest utility value. This framework can be expressed as:

$$util(x_i, f_p) \geq util(x_j, f_p) \forall j \implies i \succ j \forall j \in C$$

where  $util()$  is the utility function,  $x_i, x_j$  are vectors of attributes describing mode alternatives  $i$  and  $j$  (e.g, travel time, travel cost, etc.),  $f_p$  is a vector of characteristics of the decision maker that influence their preferences (e.g., income, number of automobiles),  $i \succ j$  means that  $i$  is preferred to  $j$ , and  $C$  is the choice set of various alternatives.

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:

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

$\gamma_k$  is the parameter which defines the direction and importance of the attribute  $k$  on the utility of an alternative  $i$ . Similarly,  $\lambda_l$  defines the direction and importance of the person's characteristic  $l$ .

$\epsilon(x_i, f_p)$  is the unobservable portion - the error term - of the utility function. Prior work (Koppelman & Bhat, 2006) in transportation mode choice recommends an assumption that the error terms of various alternatives are independent. This assumption leads to a *multinomial logit model*. This model estimates the probability  $Pr(i, p)$  of the person  $p$  selecting the alternative  $i$  from the set of alternatives  $C$ :

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

Empirical work in transportation research (Ben-Akiva et al., 1985) has studied how  $\gamma$  and  $\lambda$  parameters in equation 1 can be estimated from human choice data using maximum likelihood estimation given equation 2.

### 3.3 Limitations & Challenges for Influence

While each of these threads provides insights for the influence problem, there are significant gaps. While we incorporate many of the ideas from AI planning including constraints on multi-modal transit graphs and time-dependent edges, finding optimal routes depends on computing costs for traversing each edge in the transportation network which poses two

challenges in our context: (1) They include a large number of user parameters that are difficult to elicit from users (e.g., how much do you value time riding a bus versus a train versus driving a car?), and (2) Empirical work from transportation choice theory indicates that there are attributes that do not correspond to edge weights (e.g., household income, number of vehicles at home). We are unable to apply choice theory directly, because by itself this theory does not suggest how mode alternatives can be generated.

Furthermore, neither of these methods explicitly provide a definition of *acceptability*, *adoption*, or influence. Planning computes the optimal path for an individual given a cost function defined for every edge and choice theory predicts the probability of a trip occurring given some observable characteristics about the person and the mode of transport. The results of these computations are not the same as an estimate of the likelihood that the user will take an alternative mode if recommended (mode adoption likelihood). This is necessary for a solution to our decision-theoretic problem. That is, 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. In the sections below, we propose how these challenges can be addressed for *acceptable* planning.

#### 4. Defining Acceptability

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

Following probabilistic choice theory (equation 2):

$$\begin{aligned}
 \frac{e^{val(x_u, f_p)}}{e^{val(x_r, f_p)}} &= \frac{Pr(u, p)}{Pr(r, p)} \\
 e^{val(x_u, f_p) - val(x_r, f_p)} &= \frac{Pr(u, p)}{Pr(r, p)} \\
 \Delta_{u,r,p} = val(x_u, f_p) - val(x_r, f_p) &= \ln \frac{Pr(u, p)}{Pr(r, p)} \\
 \Delta_{r,u,p} = -\Delta_{u,r,p} &= \ln \frac{Pr(r, p)}{Pr(u, p)}
 \end{aligned}
 \tag{3}$$

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. Exact values of coefficients are useful for comparing various attributes in the model to understand the trade-offs a person may be willing to make and for computing monetary incentives. For influence, this exact computation of utility value is not necessary and the switching gain (and acceptability) can be computed from estimated likelihoods alone.

However, before acceptability is integrated in the planning framework, the following questions merit further exploration.

- What is the exact quantity that represents acceptability? Based on the discussion here, we begin by considering three candidate definitions:
  1. switching gain:  $\Delta_{r,u,p}$
  2. odds:  $e^{\Delta_{r,u,p}} = Pr(r,p)/Pr(u,p)$
  3. likelihood of the recommended mode:  $Pr(r,p)$

A related question is - how can acceptability be computed for a person when there is no prior observed mode selection data for them?

- Does acceptability of a mode affect its adoption when it is included in a recommendation? Are there other factors that may influence mode adoption above and beyond a recommendation’s acceptability?

We answer these questions through an empirical approach described in the following sections. In Section 5, we develop a machine learning approach using a large-scale, publicly available survey dataset to estimate mode likelihood for a person taking a given trip. These likelihoods can be used to compute acceptability using any candidate definition. Then, we compare different candidate definitions of acceptability (computed using these likelihoods) in Section 6 determine which definition best predicts adoption.

## 5. Estimating Acceptability: A Machine Learning Approach

As explained earlier, in order to estimate acceptability of a recommendation  $r$  for a person  $p$ , we need to calculate the likelihood of taking the mode of the recommendation,  $Pr(r,p)$  as well as for their usual mode,  $Pr(u,p)$ . Here we explore how a data-driven, machine learning method can be used to estimate these likelihoods. A typical ML classifier, such as a random forest classifier, generates a prediction of a class, given a set of input features, by computing probability for each class label and selecting the class with the highest probability. Thus, the classifier not only predicts the class label, but also the probability for every other class label. This probability estimate is critical for our approach.

### 5.1 Dataset

To develop our ML classifiers, we used the California Household Travel Survey (CHTS; California, 2012) that consists of a single-day (distributed in a year) travel diary of people



from 58 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 belonging to 2,006 households. Each sample in this dataset is a trip (defined by an origin-destination pair) taken by a specific individual. This dataset contains 78,380 trips undertaken in Los Angeles county. For each trip, we extracted the following 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.

These trips were made using 7 different modes - walk, cycle, bus, subway/train, drive, auto passenger, and motorcycle.

## 5.2 Models

We explored two different multi-class prediction problems: mode prediction and category prediction. The mode prediction problem is predicting the most likely mode given some characteristics about the trip and the person making that trip. For category prediction, we created additional labels with walking & cycle as *non-motorized*, driving, riding, & motorcycle as *motorized*, and bus & subway as *public transport*. Consequently, the category prediction problem is predicting the category of the mode given trip and person characteristics. For each prediction problem, we developed two kinds of classifiers using standard open-source implementations:

1. Random forest (Pedregosa et al., 2011): an ensemble of decision trees with 20 estimators of 30 depth each.
2. Multi-layer perceptron (Martin Abadi et al., 2015): a multi-layer perceptron with 4 layers ([1000, 500, 100, 100]).

We selected these classifiers because they impose minimal constraints on the relationships between features and class labels and can capture non-linear decision boundaries. As a starting point of exploring the prediction problems, these types of classifiers are fairly reliable. The specific parameters were found by a brief, bucketed linear hyperparameter search.

To study if these classifiers captured useful predictive information, we compared their performance with two baselines:

1. *most frequent*: This baseline assigns the most frequent class label to every sample.
2. *weighted random*: This baseline randomly assigns class labels to each sample based on its distributional frequency in the dataset.

Mode	Baseline 1 (most frequent)	Baseline 2 (weighted random)	Random Forest	Multi-layer Perceptron
<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
<b>Total</b>	0.68	0.40	0.88*	0.74
Mode category				
<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
<b>Total</b>	0.68	0.70	0.94*	0.86

Table 1: F1-scores for predicting mode classes and major modes with various machine learning methods. \* indicate the best performing classifier for that mode.

### 5.3 Results

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. These results suggests that the ML classifiers contain useful predictive information that can be used to estimate the switching gain (as well as acceptability of a recommendation). We can see that the random forest classifier has better performance than the multi-layer perceptron.<sup>1</sup> 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. Given that the random forest classifier achieves the best performance on this dataset, this classifier was selected for implementing the remainder of our approach.

A deeper look (Figure 2) into the random forest classifiers reveals the Gini importance (Breiman, 1984) of various features in predicting the mode and category. As expected, we can see that the distance of the trip is the best predictor. People prefer to drive when they have to travel longer distances and there is no reliable public transport option. Apart from trip distance, characteristics of a person’s transportation network (such as accessibility of

1. With additional effort and design, there is room for significant improvement in each of these classifiers. The results presented were sufficient for our objective of evaluating the overall framework.

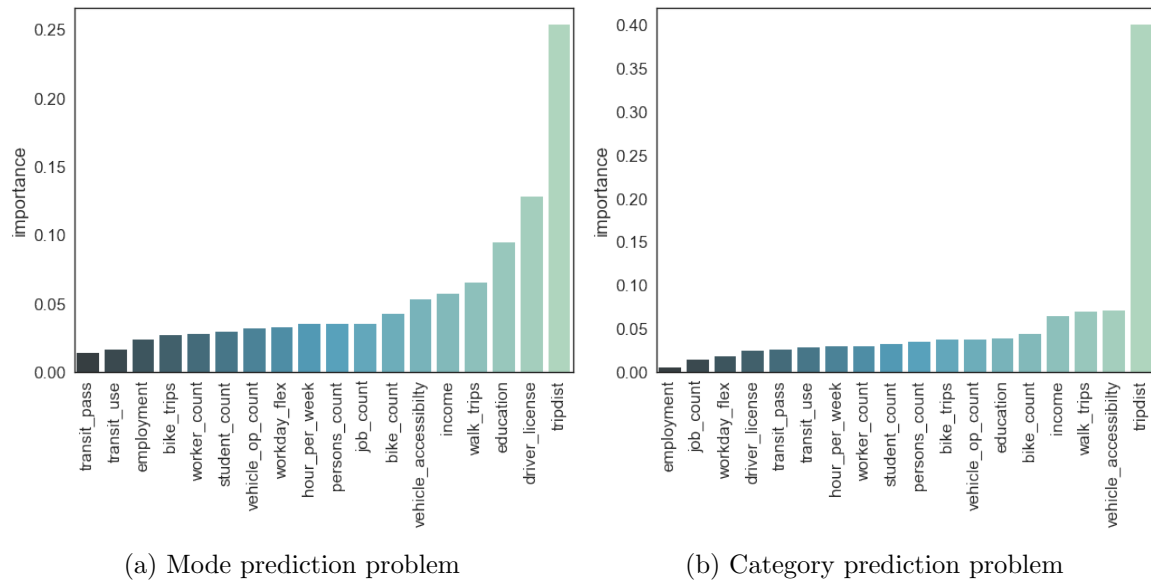


Figure 2: Gini importance (Breiman, 1984) of various features in predicting mode and category using the random forest model

vehicles in their home) and their past experience (such as walking trips undertaken last week) also contain predictive information.

### 5.4 Towards Computing Acceptability

For our candidate definitions of acceptability (in Section 4),  $P(r, p)$  and  $P(u, p)$  are important quantities. In this section, we demonstrate how these quantities can be estimated from a population survey dataset. In particular, a person  $p$  is defined as a collection of attributes encoding their demographic and employment information. Our ML approach is a significant step forward from previous empirical approaches in transportation research (Ben-Akiva et al., 1985) that have relied on standard features such as distance, income, etc. to predict mode selection. We conducted detailed interviews and survey (Mohan, Klenk, & Bellotti, 2019) and found that a large variety of factors play a role in mode selection. Our ML-based approach here allows for a breadth of features including trip features (trip distance) as well as features that represent a person’s transportation network (accessibility of various modes) and their experience (walking trips in the past week). It should be noted that while these previous approaches are more detailed in measuring the relative importance of each attribute of an individual’s preference, our approach is extensible to whatever attributes are available. As we discussed in Section 4, for a planning formulation that seeks to select from various available plans, detailed models are not necessary. An approach such as ours that provides an overall evaluation will suffice.

Computing mode acceptability using these likelihoods captures the intuition that it is easier to switch to modes taken by similar people in similar contexts. Thus, our primary assumption is that the utility function (and consequently, acceptability) can be generalized to people with similar characteristics. For example, a traveller who usually drives to a

destination lives in a neighborhood where 50% people drive, 40% use transit, and 10% walk to the same destination. Intuitively, it should be easier for the traveller to switch to transit than to walk.

## 6. Evaluating Acceptability: Evidence from a Stated Preference Study

Recall from Section 4 that not only do we have three candidate definitions of acceptability, it has not been previously shown that acceptability of a mode (computed using the ML models) will in fact impact adoption. Further, it is unclear if individuals differ in how acceptability of a recommendation impacts its adoption. These differences may arise due to difference in demographics (i.e., young individuals are more open to trying newer ways to travel) or due to some aspects of their personality that make them more or less susceptible to persuasion (Cialdini, 2001). In this section, we study these questions. We rely on stated preference methods (Johnston et al., 2017) that are used to elicit customer preference as well as for product valuation via survey responses. In particular, we conducted a small choice experiment for mode adoption. The survey instrument was adapted for our use-case of recommendations made through a travelling assistant application. This experiment had the following goals:

1. to establish that when recommendations are provided to regular drivers, they will consider adopting the recommended mode;
2. to evaluate if adoption of the recommended mode depends on its acceptability computed via the ML models in Section 5;
3. to study individual factors in the impact of acceptability on mode adoption; and finally,
4. to choose a specific definition of acceptability and develop a working model of adoption to integrate with planning.

A choice experiment presents a participant with a set of options that differ on various attributes and asks the participant to select one option. This experiment and the following analysis uncovers how participants value different features. In our context of mode recommendation, one of the options is traveling by driving - the *usual* mode - and another option is the *recommended* mode determined by COPTER through acceptable planning. We designed a choice experiment that mimics our context as described next. Our experiment had 49 (27 female, 22 male) participants who are regular drivers in the Los Angeles area. The participants had a mean age of 38.98( $\pm 12.25$ ) and a median household income of \$75,000 – \$99,999. The participants were distributed in the Los Angeles area and have diverse commuting behaviors as shown in Figure 3. The participants were gathered by a paid recruiter who ensured compliance with responses. The participants were provided Amazon gift cards worth \$30 for completing the study.

### 6.1 Materials

Each study participant took two surveys within a week:

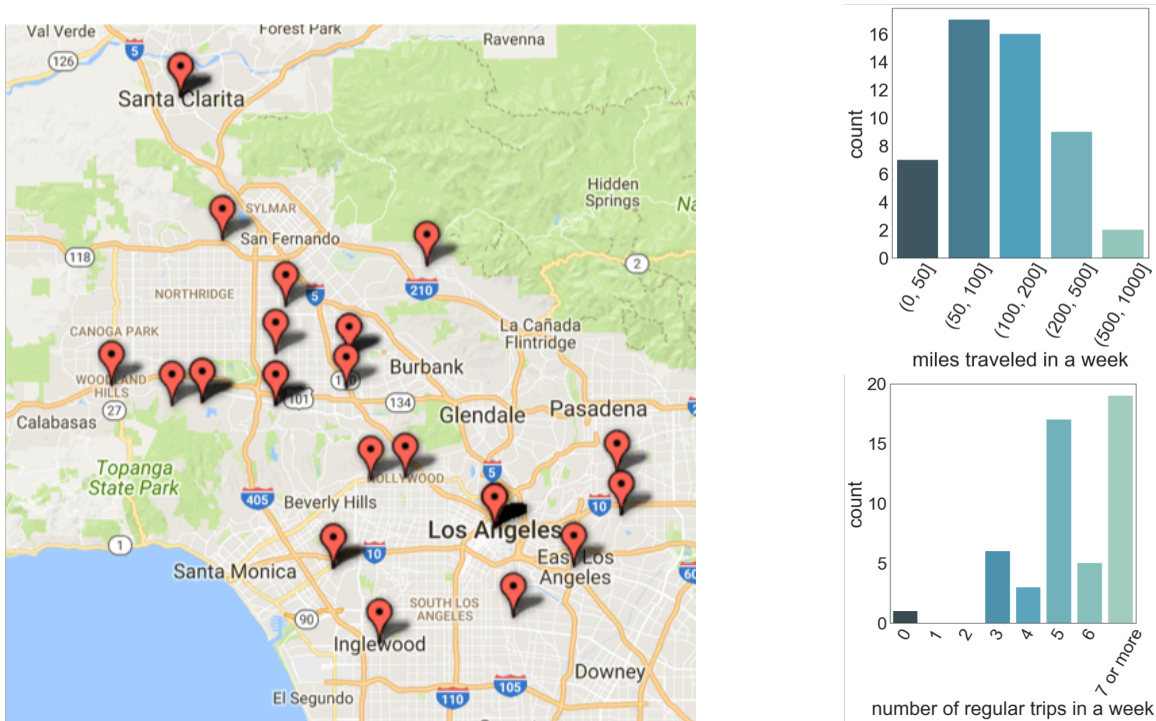


Figure 3: Location of study participants in the Los Angeles Metropolitan Area and their weekly travel characteristics.

1. *Profiler*: For each participant, this survey collected a set of trips they regularly made by driving in terms of the origin, destination, departure time, and purpose. Participants were asked to provide a specific name for each trip (such as *Film Networking Events* in Figure 4). This name was used in a later survey to refer back to this trip. Additionally, we measured participants on the features incorporated in our random forest classifier (from Section 5). To study individual differences in the impact of acceptability on adoption, we collected information about participant demographics: their age and gender as well as their susceptibility to persuasion using Cialdini’s scale (Cialdini, 2001). The 17 item scale measures a participant’s agreement with various values that underlie persuasion and influence: *reciprocity* (2 items), *commitment* (2 items), *social proof* (2 items), *authority* (2 items), *liking* (2 items), and *unity* (5 items). Measurements are done on a 7-point Likert agreement scale from *strongly disagree* to *strongly agree*.
2. *Mode adoption*: The mode adoption experiment was generated programmatically from the data collected in the profiler survey. We generated 10 mode adoption questions for each participant. For all trips reported in the profiler survey, we generated a set of available modes. This set included: *public transit*, *driving slow*, *driving 15 minutes later*, *biking* if the trip distance is < 3 miles and the participant owns a bicycle, and *walking* if the trip distance is < 1 mile. Each available mode must correspond to a class from the mode prediction problem in the previous section. We included driving-

related modes as a sanity check for the responses in study. Because both driving modes ask participants to change their behavior, and consequently, will result in a negative switching gain. We expect that due to this negative switching gain, they will not be fully adopted. Further, because both are driving-related (and are similar to what the participants are already doing), their adoption will be better than taking the public transport or walking. Evidence suggests that the driving-related modes can reduce transportation energy consumption. From the trip-mode pairs, we randomly selected 10 pairs. For each trip-mode pair, we generated a recommendation to be presented to the participant. This recommendation (see Figure 4) included information about the original trip: name of the trip, origin and destination addresses, day of the week, usual departure time, and usual arrival time; and the new recommended trip: recommended mode, departure time, and expected arrival time. Additionally, a visual representation of the proposed route was shown. The visual representation was generated using a web application - TRIPGO TripGo, 2019. The participants were provided the context of this recommendation and asked to evaluate if they will take the recommendation.

This question sets up the choice experiment. Participants had indicated that their usual mode of travel for that trip is driving. This recommendation sets up an alternate scenario and the participants are asked if they are willing to change their behavior by accepting the recommendation. The participants indicated likelihood of accepting the recommendation using a 7-point Likert scale: from *extremely unlikely* (−3) to *extremely likely* (3). If the participant is tending towards the *unlikely* end of the scale, they can be considered choosing their usual mode, driving. When the participant is tending towards *likely*, they can be considered choosing the recommended mode.

## 6.2 Data Analysis

From the materials described above, we have the following for each mode adoption measurement:

1. Independent variables
  - (a) *Computations for acceptability.* We estimated the probability of driving, taking public transit, and taking a non-motorized mode (cycling or walking) using the random forest ML model for category prediction described previously. Using these probabilities, we computed the following three values of acceptability as per the three candidate definitions of acceptability:
    - i. switching gain:  $\Delta_{r,u,p}$
    - ii. odds:  $e^{\Delta_{r,u,p}} = Pr(r,p)/Pr(u,p)$
    - iii. likelihood of the recommended mode:  $Pr(r,p)$
  - (b) *Individual variability.*
    - i. *Demographics*, measurement of gender and age.
    - ii. *Susceptibility to persuasion*, measurements on 17 item Cialdini scale.
2. Dependent response variables
  - (a) *Ordinal adoption:* Participant response on the 7-point Likert scale.



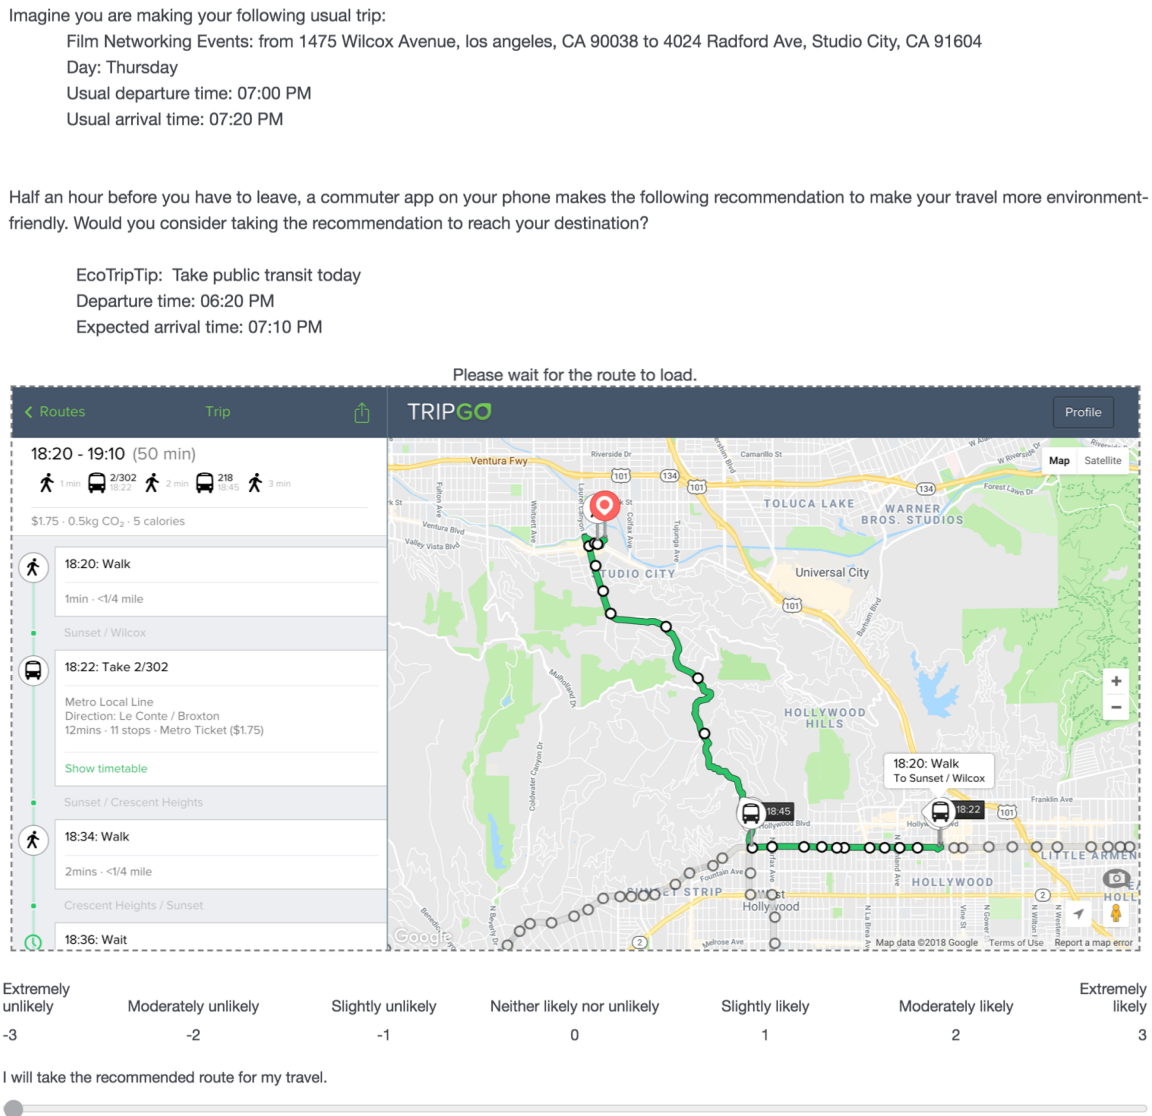


Figure 4: An example of the mode adoption question in the survey along with the response measurement instrument.

(b) *Binary adoption*: 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.

Using the data above, we conducted the following analyses:

1. In Section 6.3.1, we plot the distribution of ordinal adoption responses for various modes to show that participants consider adopting modes when provided with a recommendation and that our instrument - the mode adoption survey - is a reasonable instrument.



2. In Section 6.3.2, we evaluate if acceptability of a mode (computed using population-based ML models) is predictive of adoption and which candidate definition should be adopted. To do this, we employ mixed-effect linear and logistic regression models for the ordinal and binary adoption respectively. Mixed-effects linear regression is of the form  $y = \alpha + \beta x + \gamma z + \epsilon$ , where  $y$  is the adoption vector,  $x$  is a fixed-effect vector corresponding to different definitions of acceptability, and  $z$  is a random effect vector corresponding to a participant.  $\alpha$  is the intercept,  $\beta$  and  $\gamma$  are coefficients, and  $\epsilon$  is the error term. 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). Participants are modeled as random effects in these models under the *random intercept* assumption which encodes that some variance in responses is due to each participant having a different baseline reaction to recommendations. We report  $R^2m$ , the proportion of variance in  $y$  explained by measured, independent variables as well as  $R^2c$ , the proportion of the variance that is explained by the complete model including the random individual-level effects.
3. In Section 6.3.3, we explore the individual differences in the impact of acceptability in adoption of recommendations. Here, we include participants' measurements on the Cialdini's scale as fixed-effects in the regression models as well as their age and gender.
4. In Section 6.3.4, we select a definition of acceptability and develop a working model of adoption for formulation of acceptable planning.

## 6.3 Results

### 6.3.1 ADOPTION OF DRIVING v/s OTHER MODE RECOMMENDATIONS

First, we compare the adoption rates of driving-related modes with others. As discussed earlier,  $H1$ : we expect the adoption rates of driving-related modes to be higher than other modes as they are similar to what the participants are already doing. However,  $H2$ : we expect them to not be fully adopted as they do require a change in usual behavior.

We compare ordinal adoption of various recommended modes by plotting Likert-scale ratings as in Figure 5. We see that both  $H1$  and  $H2$  are supported in our data. Further, we find that adoption of *driving slow* may be better than that of *driving later*. That is reasonable, given that for some participants arriving at their destination by a certain time may be required. Given that the data collected through the choice experiment study matches our intuition about the domain, we can conclude that the study is measuring useful information about participants' behavior.

### 6.3.2 ACCEPTABILITY AND MODE ADOPTION

Next, we study different definitions of acceptability proposed in Section 4. Recall that we have three candidate definitions:

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

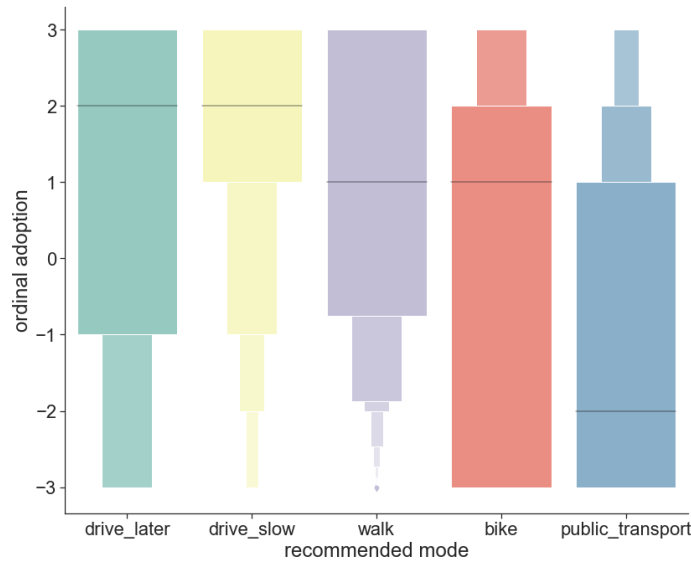


Figure 5: Box plots showing the distribution of ordinal adoption responses by mode.

where  $r$  is the recommended mode,  $u$  is the usual mode for person  $p$  and  $Pr()$  is the likelihood of prediction from our ML models. Results from mixed-effect linear and logistic regression are shown in Table 2. For each definition of acceptability, we report  $R^2c$ , the proportion of the variance that is explained by the complete model, and  $R^2m$ , the proportion is explained by the candidate definition. Each cell in the table denotes a fitted mixed-effects regression model.

Overall, the results show that people reported a willingness to change their travel behavior by adopting recommended modes. Further, we see that all definitions of acceptability significantly impact the adoption of a recommendation. The coefficients align with the intuition about the domain - as the cost of adopting the recommendation decreases (acceptability increases), adoption increases. We see that  $R^2c$ ,  $R^2m$  are highest for odds. These results suggest that it is the best predictor of mode adoption because it explains the most variance in the data. The odds and probability have similar impact on adoption due to the fact that our participants are regular drivers and  $Pr(u, p) \approx 1$  for many of the trips. With a more diverse population in a larger deployment, the differences between these definitions of acceptability can be teased apart. For further analyses and implementation, we chose the odds definition of mode acceptability.

Figure 6 shows the distribution of responses and the results of the regression model with odds as the independent variable. The green points represent adoption choices by different participants. The black line shows the cumulative density of responses. The blue line and ribbon capture the coefficient of odds and the confidence interval. From the graph, it can be seen that the statistical significance arises from the responses where the estimated likelihood ratio  $> 0.5$  ( $\leq 10\%$  of responses). For lower values of likelihood ratio, the responses are evenly distributed. One explanation of these observations is that when asked about changing behavior to support sustainable modes of transport, participants tended to be optimistic. In a deployment where real choice data is available, revealed preference

Dependent → Independent ↓	Adoption (ordinal)	Adoption (binary)
(intercept)	-0.017	-0.185
switching gain, $\Delta_{r,u,p}$	0.108*	0.104*
$R^2_m$	0.034	0.035
$R^2_c$	0.347	0.270
(intercept)	-0.025	-1.065
odds, $e^{\Delta_{r,u,p}}$	2.386***	1.780**
$R^2_m$	0.075	0.064
$R^2_c$	0.379	0.346
(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. \*\*\*  $p < 0.001$ , \*\*  $p < 0.05$ , \*  $p < 0.1$




Figure 6: Participant mode adoption (green), cumulative density of responses (black), and results of the mixed-effects logit model with odds as independent variable (blue)

methods (Carson, Flores, Martin, & Wright, 1996) can be employed for more accurate choice modeling.

### 6.3.3 INDIVIDUAL VARIABILITY IN ADOPTION

Results in the previous section provide strong evidence that acceptability of a mode is a significant predictor of its adoption. Next, we study if there are individual differences in adopting a mode recommendation above and beyond its computed acceptability. We have two kinds of factors that may bring individual variability in adoption: demographics and personality - in particular, susceptibility to persuasion. We analyze the impact of these factors by extending the regression models to include these factors as independent variables in two different efforts. To be valuable, these factors must influence adoption significantly as well as increase the variance explained by fixed-effects of the model ( $R_m^2$ ).

First, in Table 3, we show results from the regressions that include an ordinal variable representing participant age and a categorical gender variable. As is apparent, there is no significant impact of either age or gender, suggesting that these variables do not play a role in adoption.

Dependent → Independent ↓	Adoption (ordinal)	Adoption (binary)
(intercept)	-1.63736	-1.67021
odds, $e^{\Delta_{r,u,p}}$	2.38833***	2.15086**
GenderMale	0.33149	0.01281
Age	0.01377	0.21983
$R^2_m$	0.085	0.072
$R^2_c$	0.396	0.310

Table 3: Regression models with age and gender for ordinal and binary mode adoption. \*\*\*  $p < 0.001$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Next, we extended the mixed-effects regression model to include independent variables from the susceptibility to persuasion scale (Cialdini, 2001). In Table 4, we show the results from these regressions. As previously, we model both ordinal and binary adoptions, shown as models of class I and II in the table. For each class, we develop three models with varying independent variables. Models IA & IIA are reproduced from Table 2 for comparison. Models IB & IIB are full models and include all measurements from the Cialdini scale as independent variables. Models IC & IIC are simplified models that only have variables judged to be significant in the full models.

In models IB, we see that commitment (i7) and liking (i12) dimensions are significant predictors of adoption. As is expected,  $R_m^2$  and  $R_c^2$  both increase significantly suggesting that including these variables in the model helps explain more variance in our dataset, above and beyond what is explained by odds. In model IC, in which we removed non-significant persuasion items, we see that both commitment and liking dimensions are still

MODEL	I			II		
	A	B	C	A	B	C
Dependent →	Ordinal adoption			Binary adoption		
Independent ↓						
(intercept)	-0.025	-0.183	-0.201	-1.065	-0.084	-0.998
odds, $e^{\Delta_{r,u,p}}$	2.386***	2.306***	2.348***	2.1587**	1.780**	1.729**
I1 <i>Reciprocity repay favor</i>		-0.138			-0.166	
I2 <i>Reciprocity return gift</i>		-0.009			-0.043	
I3 <i>Scarcity last to buy</i>		-0.034			-0.009	
I4 <i>Scarcity special value</i>		-0.212			-0.170	
I5 <i>Authority obey directions</i>		0.311			0.330*	0.1386
I6 <i>Authority listen to</i>		-0.002			0.021	
I7 <i>Commitment appointment</i>		0.495*	0.515**		1.126**	0.979***
I8 <i>Commitment promise</i>		0.448			0.207	
I9 <i>Consensus advice</i>		-0.051			0.031	
I10 <i>Consensus what to do</i>		0.170			0.206	
I11 <i>Liking inclined to believe</i>		-0.069			0.075	
I12 <i>Liking do favor</i>		-0.966**	-0.819**		-1.105**	-0.905**
I13 <i>Unity help</i>		-0.143			-0.238	
I14 <i>Unity self assembly furniture</i>		-0.026			-0.008	
I15 <i>Unity country</i>		0.236			0.128	
I16 <i>Unity sports team</i>		0.157			0.162	
I17 <i>Unity vicarious success happy</i>		-0.357			-0.491	
$R^2_m$	0.075	0.224	0.156	0.064	0.329	0.254
$R^2_c$	0.379	0.467	0.388	0.301	0.346	0.318

Table 4: Regression modeling results for ordinal and binary mode adoption. \*\*\*  $p < 0.001$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

significant. Further, we see that  $R_c^2$  in models IC is close to that of model IA suggesting that they explain similar amounts of variance. However, in model IC  $R_m^2$  (variance explained by fixed-effects) is much higher than that of model IA. A proportion of variance that was explained by random effects (corresponding to participants) in model IA can be explained by measured persuasion items in model IC. Similar observations can be made of Models II. This suggests that persuasion items measure important predictive information about individual adoption of recommendations above and beyond what odds does.

#### 6.3.4 WORKING MODEL OF ADOPTION

In this section, we first demonstrated that providing recommendations to people in a travel route choice task can indeed influence their decisions. Then, we evaluated if acceptability (odds, as defined in Section 4 and computed using ML model in Section 5) impacts adoption. Our analysis suggests that it indeed is a significant factor in recommendation adoption. Further, we show that different individuals may be influenced to a different extent. In our analysis, a person's measurement of susceptibility to persuasion on Cialdini's scale is useful in explaining this individual variation in the impact of acceptability on 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* among other things to motivate behavior change. At this point, we have defined, computed, and evaluated what acceptability is. Therefore for the remainder of this paper, as a working model of adoption, we use the logit model of odds to estimate the probability of a person adopting a recommended route. Understandably, this is a limited model of adoption and our future efforts will extend this model further to include other factors as well. In particular, our results in Section 6.3.3 suggest that the framing of a recommendation can be personalized to every individual such that it is compelling to them. We will extend our adoption model to include the impact of personalized compelling framing.

## 7. Finding Energy-Efficient Acceptable Plans

Until now, we have defined plan acceptability, proposed a data-driven way to compute it, and have validated it to be an important component of mode adoption. Now, we extend the multi-modal planning formulation to include acceptability. The main concern in incorporating acceptability in planning is that costs in planning are defined at edges, while acceptability is an overall evaluation of a plan recommendation. For integration, we sacrifice efficiency and split plan evaluation in two phases as described below.

Figure 7 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 with the following steps:

1. Generate a mode candidate set  $M_p$  for the person  $p$  for whom the request is made.  $M_p$  can be generated by knowing 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  $L_p$  set that is valid. For the example above,  $L_p = \{w*, w * b + w*, w * s + w*\}$ . 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. This process will generate a candidate set of plans for the person,  $\Pi_p$ .
4. Compute the energy reduction in each plan  $\pi \in \Pi_p$  using an existing mesoscopic energy model (Elbery, Dvorak, Du, Rakha, & Klenk, 2018).
5. Evaluate the likelihood of adoption,  $adopt(\pi, p)$ , for every  $\pi \in \Pi_p$  using the logit model of adoption in the previous section.
6. Select a plan that has maximal expected energy savings,  $\pi^* = \arg \max_{\pi \in \Pi_p} adopt(\pi, p) \times energy(\pi)$ .

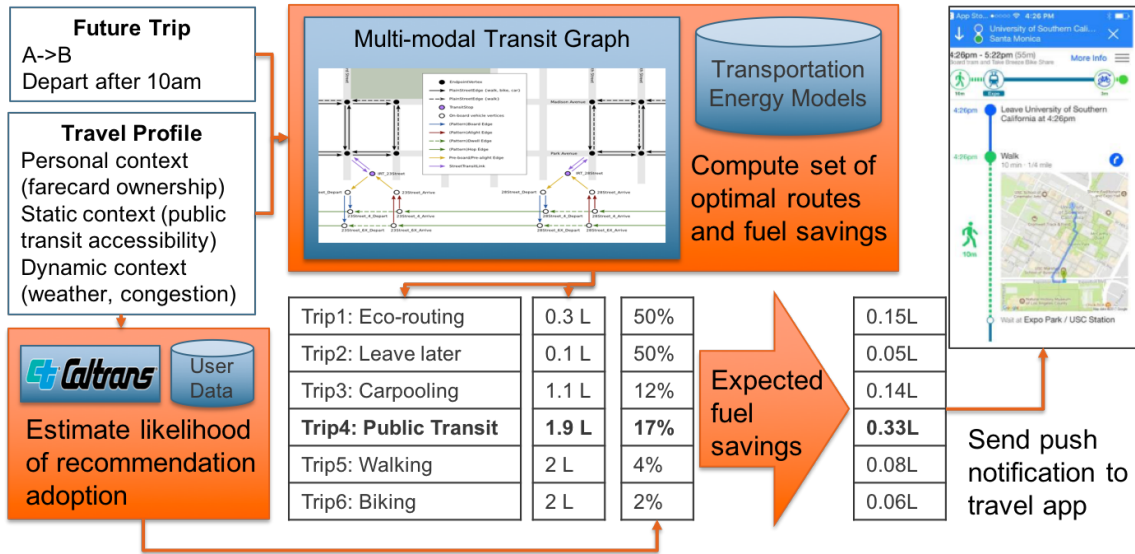


Figure 7: COPTER identifies the recommendation with the largest expected reduction in energy consumption by acceptable planning and energy modeling.

There are two items in the above process that merit additional discussion. First, is the proposed acceptable planning framework extensible to additional multi-modal route recommendations? For each element in the candidate mode set  $M_p$ , there must be a language  $L_p$  that constrains the possible sequences of modes and corresponding likelihood of adoption,  $adopt(\pi, p)$ . In the empirical work discussed in this paper, the likelihood of adoption is the same for all trips involving public transit (an element of  $M_p$ ). To consider trips where the traveler walks to transit distinctly from those in which the traveler takes a taxi for the first/last mile requires changing  $M_p$  and  $L_p$  as well as gathering mode adoption data to update  $adopt(\pi, p)$  for the new modes in  $M_p$ . Computing  $adopt(\pi, p)$  requires the acceptability of the plan which in turn is estimated from travel behaviors of a population. As data collection in cities gets more sophisticated, we expect to have a richer dataset of mode sequences for trips. This enables the use of classification-based likelihood estimation to compute adoption.

Second, to clarify the plan selection (Step 6), consider two options for a particular traveler: walking which saves huge amounts of energy but its likelihood of adoption is very low and taking public transit which does not save as much energy but has much higher likelihood of adoption. Based on quantitative expected savings in fuel, COPTER may recommend walking by ranking it above public transit, which may not result in adoption and consequently, no savings in fuel. If plan selection is biased by likelihood of adoption alone, COPTER will recommend public transit which has improved chances of adoption and consequently, is more likely to result in some energy savings. From a behavior change perspective, the second approach may be better - some change is better than none at all. This is a design question that can be answered empirically by deploying both competing plan selection approaches. However, in absence of empirical observations from deployments,






Figure 8: Los Angeles region study area.

we consider the algorithm above for further evaluation. Our formulation allows for exploring both design criterion in future.

## 8. Energy Simulation Experiments

While deploying COPTER in a large city to measure the impact is the ideal evaluation, it is extremely resource intensive to conduct. Furthermore, without being incorporated into a widely used product, it would be impossible to iteratively evaluate different components of a complex system such as COPTER employs. Therefore, we evaluate the potential impact of COPTER using a high fidelity simulation of the Los Angeles region as shown in Figure 8. Our integrative simulation simulates people of Los Angeles making commuter trips in the morning and evening periods. It uses the adoption model developed in the previous sections to simulate COPTER’s influence on their decisions and computes energy impact using a state-of-art transportation energy simulation of the LA region described briefly below.

### 8.1 Transportation Energy Model

The simulation model (Elbery et al., 2018) 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, Elbery, Rakha, & Klenk, 2018). The simulation incorporates previously validated driver behavior (Rakha, Pasumarthy, & Adjerid, 2009) and fuel (Fiori, Ahn, & Rakha, 2016) models to provide accurate estimates of travel time and energy consumption due to traveler influence.

## 8.2 Simulation Scenario

Our study area, shown in Figure 8, covers the majority of Los Angeles County. 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), our calibrated transportation model includes 1.2 and 1.7 vehicle trips respectively. As the effects of congestion continue past the peak period, our study considers an AM peak/off-peak period (7am-12pm) and a PM peak/off-peak period (4pm-9pm).

In our deployment scenario, we assume that 10% of peak period drivers use apps (e.g., Waze, TripGo) through which we can offer alternative suggestions.

## 8.3 Simulation Method

Given our deployment assumption that 10% of the driving population receives our messages, we sample 10% of the traffic during the peak periods (7am-10am and 4pm-7pm) for our influenced population. Our experiment has two conditions:

1. Baseline: Each influenced traveler drives their original route
2. Influence: Each influenced traveler probabilistically takes the maximum expected energy reduction route as determined by COPTER using the mesoscopic model or drives their original route based on working model of adoption in Section 6.3.4.

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, 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.

For simulation in the influence condition, we:

1. Sampled features from the CHTS dataset and used these features to augment trips in our energy models
2. For each trip, selected a recommendation from motorized, non-motorized, and public transit using COPTER
3. Computed the likelihood of adoption by applying our working model of adoption as follows:
  - (a) Generated an intercept by fitting and sampling from a Gaussian distribution with intercept mean and standard deviation from our analysis.

(b) Computed acceptance likelihood by applying the working model of adoption

4. Simulated adoption by selecting the recommendation biased by likelihood of adoption

### 8.4 Results

Table 5 contain the results for AM and PM periods obtained using the microscopic model. 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 6 indicates that almost half the influenced population is willing to take alternative modes. These mode switches are based on our mode adoption study and may be optimistic. Future work will refine these estimates through a deployment of the technology and measuring people’s behavior.

	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%)

(a) Results of departures in the AM period: 7:00 - 12:00

	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%)

(b) Results of departures in the PM period: 16:00 - 21:00

Table 5: 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.

## 9. Related Work

This paper incorporates ideas from diverse disciplines and consequently, is related to a large body of work in transportation and AI. Here, we briefly discuss how our work relates to ongoing research.

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

## 9.1 Transportation

The majority of transportation policy and social science research either focuses on systemic changes that are applied evenly to the entire population such as congestion pricing/dynamic tolling (Brownstone, Ghosh, Golob, Kazimi, & Van Amelsfort, 2003) or studies long term behavior change in individuals to encourage sustainable commuting (Castellani, Colombino, Grasso, & Mazzega, 2016). Recently, there has been a shift toward personalizing the interventions to increase the use of sustainable transportation. IncenTrip (Sun & 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 the allocation comes from a points-based system in which companies give travelers points for alternative transportation choices that can be redeemed at local businesses. While these systems primarily rely on monetary incentives to affect travel behavior, Metropia (Zhu, Hu, & Chiu, 2018) has shown that congestion information alone can influence traveler departure time. Our work differs from the above by identifying a single personalized recommendation from diverse alternatives based on network conditions without any monetary incentives.

In addition to algorithmic work in route planning, the AI community has explored other aspects of the transportation influence problem. Work on compliant micro-tolling assumes a subset of agents would be willing to opt-in to a congestion tolling regime (Hanna, Sharon, Boyles, & Stone, 2019). In that work, energy and traffic are modeled mesoscopically to compute marginal cost paths for social benefit and human behavior is captured using a single parameter for value of time. Recent work on recommender systems use more complicated user models to order potential routes to a user (H. Liu et al., 2019). Our work differs by computing alternative routes to an expected default.

Other work in AI seeks to make predictions about the transportation network. Using position traces and other data, systems predict the mode and destination of travelers (Liao, Patterson, Fox, & Kautz, 2007; Song, Kanasugi, & Shibasaki, 2016). These works address travel prediction problems, which are essential to the commercial applicability of our approach. With respect to the transportation influence problem, COPTER could incorporate these approaches as follows. First, predicting expected future trips is necessary to trigger the generation of alternatives. Second, the prediction of current mode supports monitoring

travelers to determine if they follow the recommendation and to assist them as conditions change.

## 9.2 AI for Social Good

Hager et al. (2017) identified urban computing and sustainability as one of the motivating domains for research on intelligent methods for public welfare. The methods developed in this paper take a unique perspective towards the problem of urban computing and sustainability in recognizing that humans are a critical component of urban transportation systems. It puts reasoning about human behavior as the center of intelligent system design. Our approach aligns with the perspective of Abebe and Goldner (2018), who argue that AI for social good inherently is an inter-disciplinary problem and developing good AI solutions requires deep grounding in the field of interest and collaboration with domain experts. To develop our AI planning-based approach, we engage with behavioral economics to understand how people choose and with transportation research to estimate energy consumption.

Our work is a part of growing body of work in AI that seeks to achieve social good goals. Intelligent systems have been employed for allocation of security resources at ports (Shieh et al., 2012), protecting biodiversity in conservation areas (Fang et al., 2016), and screening 800 million airport passengers annually throughout the USA (Brown, Sinha, Schlenker, & Tambe, 2016). Research has also begun exploring intelligent algorithms and systems that can influence people’s behavior to benefit society. Most of this work has been focused on health. Hartzler et al. (2016), Mohan, Venkatakrishnan, and Hartzler (2019), Mohan, Venkatakrishnan, Silva, and Pirolli (2017) develop an adaptive interactive system for scaffolding learning of exercising behaviors in sedentary populations. Yadav et al. (2018) explore social network models to maximize influence for HIV awareness in homeless youth population. Our work brings technology-mediated influence and behavior change to the domain of urban computing and sustainability.

## 9.3 Human-Aware AI Systems

AI research on modeling of humans in human-AI collaborative systems is in its infancy. While the community has recognized the importance of incorporating explicit models of human behavior in AI systems to make them more effective in collaborative tasks (Albrecht & Stone, 2018; Kambhampati, 2019), there is a dearth of computational models that can be incorporated with AI algorithms. In problems where the humans (or their teams) are adversaries to the AI algorithm, various game-theoretic modeling techniques (Sinha, Fang, An, Kiekintveld, & Tambe, 2018) have been found to be useful. In dearth of detailed models of why a specific human behavior occurs, these approaches optimize against the worst case, providing guarantees of reasonable performance in the expected case. Often in collaborative modeling, task models employed by AI systems for acting are used as a stand-in for representing human goals and behaviors (Chakraborti, Sreedharan, Kulkarni, & Kambhampati, 2018; Kim, Chacha, & Shah, 2015). These approaches are useful when humans and systems are engaged in the same task and have similar goals.

Our proposed method takes a unique perspective towards building models of human collaborators. In our problem, the human and the agent are engaged in different tasks. The human’s goal is to reach their destination while the agent’s goal is to perturb human

cognitive system such that the human takes a sustainable mode to get there. We adopt the choice modeling techniques from behavioral economics and explore how it can be incorporated in the AI planning framework to produce personalized plan recommendations. In contrast with adversarial modeling techniques, our work incorporates explicit causal models of human decisions to produce effective solutions. Our work can be considered as an instance in the human-aware AI design agenda that exploits detailed, explicit models of human decisions, learning, and behavior and of which cognitive tutors (Koedinger, Brunskill, Baker, McLaughlin, & Stamper, 2013) are another instance. Along these lines, Mao and Gratch (2012) proposed human modeling for AI systems that is grounded in theory of human cognition and reasoning. It is worth noting that this work aims at only predicting human behavior while our work takes a step beyond prediction and employs predictive information for causing a change in behavior.

## 10. 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 evaluate different definitions of acceptability using a human mode choice study. Finally, in a simulation experiment over the Los Angeles area we show that if 10% of 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.

In future work, we would like to integrate these results with 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 and measure the switching costs associated with carpooling to expand the set of alternatives.

From a social good perspective, COPTER is ready for deployment with influenced messages sent through the TripGo application (TripGo, 2019). The next step is to evaluate the influence model by measuring actual traveler behavior and then feeding those results back into our simulation study. In a deployed scenario, we will be able to integrate additional transportation influence features identified from our survey (Mohan, Klenk, & Bellotti, 2019) (e.g., weather) by updating the acceptability model, if we have enough historical trip behavior, or the influence model by incorporating these features as parameters in the adoption model. Later work will define and assess the impacts of timeliness of a recommendation and compelling messaging by considering Cialdini's influence strategies (Cialdini, 2001).

Urban transportation is a huge problem for the environment and citizens' quality of life. COPTER demonstrates how multiple AI techniques can be adapted and combined to alleviate congestion and improve the environment through transportation influence.

## 11. Acknowledgments

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 Filip Dvorak and Aaron Ang for engineering the system, Frances Yan and Victoria Bellotti for contributions to the design and execution of the human studies presented here, and Spenser Anderson for reviewing a draft of this work. The authors appreciate the effort invested by the City of Los Angeles and CalTrans in collecting the data necessary for developing the. Finally, the authors would like to thank the anonymous reviewers for the AAAI/AIES conference on AI, Ethics, and Society and the Journal of AI Research for providing invaluable feedback that has made the writing impactful.

## References

- Abebe, R., & Goldner, K. (2018). Mechanism Design for Social Good. *AI Matters*, 4(3), 27–34.
- Albrecht, S. V., & Stone, P. (2018). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258, 66–95.
- Azevedo, C. L., Seshadri, R., Gao, S., Atasoy, B., Akkinepally, A. P., Christofa, E., . . . 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., . . . Werneck, R. F. (2016). Route Planning in Transportation Networks. In *Algorithm Engineering* (pp. 19–80). Springer International Publishing.
- Ben-Akiva, M. E., Lerman, S. R., & Lerman, S. R. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT press.
- Benoit, W. L., & Benoit, P. J. (2008). *Persuasive messages: The process of influence*. Blackwell Malden, MA.
- Botea, A. (2016). Hedging the Risk of Delays in Multimodal Journey Planning. *AI Magazine*, 37(4).
- Breiman, L. (1984). *Classification and regression trees*. Routledge.
- Brown, M., Sinha, A., Schlenker, A., & Tambe, M. (2016). One Size Does Not Fit All: A Game-Theoretic Approach for Dynamically and Effectively Screening for Threats. In *AAAI* (pp. 425–431).
- Brownstone, D., Ghosh, A., Golob, T. F., Kazimi, C., & 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.



- California. (2012). California Household Travel Survey. <https://dot.ca.gov/programs/transportation-planning/economics-data-management/transportation-economics/ca-household-travel-survey>. Accessed: 2019-09-08.
- Carson, R. T., Flores, N. E., Martin, K. M., & Wright, J. L. (1996). Contingent Valuation and Revealed Preference Methodologies: Comparing the Estimates for Quasi-Public Goods. *Land economics*, 80–99.
- Castellani, S., Colombino, T., Grasso, A., & Mazzega, M. (2016). Understanding Commuting to Accompany Work Organisations’ and Employees’ Behaviour Change. In *Smart Cities Conference (ISC2), 2016 IEEE International* (pp. 1–6). IEEE.
- Chakraborti, T., Sreedharan, S., Kulkarni, A., & Kambhampati, S. (2018). Projection-Aware Task Planning and Execution for Human-in-the-Loop Operation of Robots in a Mixed-Reality Workspace. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 4476–4482). IEEE.
- Cialdini, R. B. (2001). Harnessing the science of persuasion. *Harvard Business Review*, 79(9), 72–81.
- Delling, D., Dibbelt, J., Pajor, T., Wagner, D., & Werneck, R. F. (2013). Computing multi-modal journeys in practice. In *International Symposium on Experimental Algorithms* (pp. 260–271). Springer.
- Dibbelt, J., Pajor, T., & 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 Connexion with Graphs. *Numerische Mathematik*, 1(1), 269–271.
- Domencich, T. A., & McFadden, D. (1975). *Urban Travel Demand-A Behavioral Analysis*. Transport and Road Research Laboratory.
- Du, J., Elbery, A., Rakha, H., & 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*.
- Dvorak, F., Mohan, S., Bellotti, V., & Klenk, M. (2018). Collaborative Optimization and Planning for Transportation Energy Reduction. In *Workshop on Distributed and Multi-Agent Planning, ICAPS Workshops*, Delft, The Netherlands.
- EIA. (2016). U.S. Energy Information Administration, Independent Statistics and Analysis. [www.eia.gov](http://www.eia.gov). Accessed: 2016-09-30.
- Elbery, A., Dvorak, F., Du, J., Rakha, H., & 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* (pp. 103–112).

- Fang, F., Nguyen, T. H., Pickles, R., Lam, W. Y., Clements, G. R., An, B., ... Lemieux, A., et al. (2016). Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *AAAI* (pp. 3966–3973).
- Fiori, C., Ahn, K., & Rakha, H. A. (2016). Power-based electric vehicle energy consumption model: Model development and validation. *Applied Energy*, *168*, 257–268.
- Hager, G. D., Drobnis, A., Fang, F., Ghani, R., Greenwald, A., Lyons, T., ... Tambe, M. (2017). *Artificial Intelligence for Social Good*. The Computing Community Consortium.
- Hanna, J. P., Sharon, G., Boyles, S. D., & Stone, P. (2019). Selecting COmplian Agents for Opt-in Micro-Tolling. *Processings of the 33rd Conference on Artificial Intelligence (AAAI)*.
- Hartzler, A. L., Venkatakrisnan, A., Mohan, S., Silva, M., Lozano, P., Ralston, J., ... Pirolli, P. (2016). Acceptability of a Team-based Mobile Health (mHealth) Application for Lifestyle Self-management in Individuals with Chronic Illnesses. In *2016 IEEE 38th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC)*, (pp. 3277–3281). IEEE.
- Johnston, R. J., Boyle, K. J., Adamowicz, W., Bennett, J., Brouwer, R., Cameron, T. A., ... Scarpa, R., et al. (2017). Contemporary Guidance for Stated Preference Studies. *Journal of the Association of Environmental and Resource Economists*, *4*(2), 319–405.
- Kambhampati, S. (2019). Challenges of Human-Aware AI Systems. AAAI 2018 Presidential Address. arXiv: [1910.07089](https://arxiv.org/abs/1910.07089) [cs.AI]
- Kim, B., Chacha, C. M., & Shah, J. A. (2015). Inferring team task plans from human meetings: A generative modeling approach with logic-based prior. *Journal of Artificial Intelligence Research*, *52*, 361–398.
- Koedinger, K. R., Brunskill, E., Baker, R. S., McLaughlin, E. A., & Stamper, J. (2013). New Potentials for Data-Driven Intelligent Tutoring System Development and Optimization. *AI Magazine*, *34*(3), 27.
- Koppelman, F. S., & Bhat, C. (2006). A self instructing course in mode choice modeling: multinomial and nested logit models.
- Liao, L., Patterson, D. J., Fox, D., & Kautz, H. (2007). Learning and Inferring Transportation Routines. *Artificial Intelligence*, *171*(5), 311–331.
- Liu, H., Li, T., Hu, R., Fu, Y., Gu, J., & Xiong, H. (2019). Joint Representation Learning for Multi-Modal Transportation Recommendation. *AAAI*.
- Liu, X., Fritz, C., & Klenk, M. (2018). On Extensibility and Personalizability of Multi-Modal Trip Planning. In *Proceedings of the 11th Multidisciplinary Workshop on Advances in Preference Handling*, New Orleans, LA USA.
- Mao, W., & Gratch, J. (2012). Modeling Social Causality and Responsibility Judgment in Multi-agent Interactions. *Journal of Artificial Intelligence Research*, *44*, 223–273.

- Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, . . . Xiaoqiang Zheng. (2015). TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org. Retrieved from [https://www.tensorflow.org/api\\_docs/python/tf/contrib/learn/DNNClassifier](https://www.tensorflow.org/api_docs/python/tf/contrib/learn/DNNClassifier)
- Mohan, S., Klenk, M., & Bellotti, V. (2019). Exploring How to Personalize Travel Mode Recommendations For Urban Transportation. In *The Annual Workshop on Theory-Informed User Modeling for Tailoring and Personalizing Interfaces (HUMANIZE) in conjunction with ACM Intelligent User Interfaces*. ACM.
- Mohan, S., Venkatakrisnan, A., & Hartzler, A. (2019). Designing an AI Health Coach and Studying its Utility in Promoting Regular Aerobic Exercise. *to appear in ACM Transactions on Interactive Intelligent Systems*. arXiv: 1910.04836 [cs.HC]
- Mohan, S., Venkatakrisnan, A., Silva, M., & Pirolli, P. (2017). On Designing a Social Coach to Promote Regular Aerobic Exercise. In *Proceedings of the Twenty-Ninth Conference on Innovative Applications of AI at AAAI* (pp. 4721–4727).
- Nahum-Shani, I., Hekler, E. B., & Spruijt-Metz, D. (2015). Building health behavior models to guide the development of just-in-time adaptive interventions: A pragmatic framework. *Health Psychology, 34*(S), 1209.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., . . . Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830. Retrieved from <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>
- Rakha, H., Pasumarthy, P., & Adjerid, S. (2009). A Simplified Behavioral Vehicle Longitudinal Motion Model. *Transportation Letters, 1*(2), 95–110.
- Schrank, D., Eisele, B., Lomax, T., & Bak, J. (2015). 2015 Urban Mobility Scorecard. *Annual Urban Mobility Scorecard*.
- Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., . . . Meyer, G. (2012). Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1* (pp. 13–20). International Foundation for Autonomous Agents and Multiagent Systems.
- Sinha, A., Fang, F., An, B., Kiekintveld, C., & Tambe, M. (2018). Stackelberg Security Games: Looking Beyond a Decade of Success. In *IJCAI* (pp. 5494–5501).
- Song, X., Kanasugi, H., & Shibusaki, R. (2016). DeepTransport: Prediction and Simulation of Human Mobility and Transportation Mode at a Citywide Level. In *Proceedings of the International Joint Conference on Artificial Intelligence* (Vol. 16, pp. 2618–2624).
- Sun, Y., & Zhang, L. (2018). Microeconomic Model for Designing Public Transit Incentive Programs. *Transportation Research Record*.
- TripGo. (2019). TripGo. <https://www.tripgo.com/>. Accessed: 2019-09-08.

- Tversky, A., & Kahneman, D. (1986). Rational Choice and the Framing of Decisions. *Journal of Business*, S251–S278.
- Yadav, A., Wilder, B., Rice, E., Petering, R., Craddock, J., Yoshioka-Maxwell, A., . . . Woo, D. (2018). Bridging the Gap Between Theory and Practice in Influence Maximization: Raising Awareness about HIV among Homeless Youth. In *IJCAI* (pp. 5399–5403).
- Zhu, X., Hu, X., & Chiu, Y.-C. (2018). Will Information and Incentive Affect Traveler's Day-to-Day Departure Time Decisions? An Empirical Study of the Decision Making Evolution Process. In *the 97th Annual Meeting of Transportation Research Board*.