

Prosociality in Microtransit

Divya Sundaresan

*Department of Computer Science
NC State University
Raleigh, NC, USA*

DSUNDAR3@NCSU.EDU

Akhira Watson

*Department of Psychology
NC State University
Raleigh, NC, USA*

AWATSON9@NCSU.EDU

Eleni Bardaka

*Department of Civil, Construction, and Environmental Engineering
NC State University
Raleigh, NC, USA*

EBARDAK@NCSU.EDU

Crystal Chen Lee

*Department of Teacher Education and Learning Sciences
NC State University
Raleigh, NC, USA*

CCHEN32@NCSU.EDU

Christopher B. Mayhorn

*Department of Psychology
NC State University
Raleigh, NC, USA*

CBMAYHOR@NCSU.EDU

Munindar P. Singh

*Department of Computer Science
NC State University
Raleigh, NC, USA*

MPSINGH@NCSU.EDU

Abstract

We study (*public*) *microtransit*, a type of transportation service wherein a municipality offers point-to-point rides to residents, for a fixed, nominal fare. Microtransit exemplifies practical resource allocation problems that are often over-constrained in that not all ride requests (pickup and dropoff locations at specified times) can be satisfied or satisfied only by violating soft goals such as sustainability, and where economic signals (e.g., surge pricing) are not applicable—they would lead to unethical outcomes by effectively coercing poor people.

We posit that instead of taking rider preferences as fixed, shaping them prosocially will lead to improved societal outcomes. *Prosociality* refers to an attitude or behavior that is intended to benefit others. This paper demonstrates a computational approach to prosociality in the context of a (*public*) *microtransit* service for disadvantaged riders. Prosociality appears as a willingness to adjust one's pickup and dropoff times and locations to accommodate the schedules of others and to enable sharing rides (which increases the number of riders served with the same resources).

This paper describes an interdisciplinary study of prosociality in microtransit between a transportation researcher, psychologists, a social scientist, and AI researchers. Our contributions are these: (1) empirical support for the viability of prosociality in microtransit (and constraints on it) through interviews with drivers and focus groups of riders; (2) a prototype mobile app demonstrating how our prosocial intervention can be combined with the transportation backend; (3) a reinforcement learning approach to model a rider and determine the best interventions to persuade that rider toward prosociality; and (4) a cognitive model of rider personas to enable evaluation of alternative interventions.

1. Introduction

Transportation is essential for residents to go to work, obtain healthcare, shop for food, or otherwise engage in civic life. Public transportation is an important municipal service because it is essential for disadvantaged individuals, such as those who cannot afford to own a car, and is often environmentally preferable to personal transportation. Increasingly, small municipalities are finding that fixed-route bus services are expensive and underutilized and do not serve their residents well. (*Public*) *microtransit* services are a kind of public transit system wherein a municipality offers point-to-point rides to residents (Shaheen, Cohen, Chan, & Bansal, 2020). In the settings of interest, microtransit can involve a mixture of on-demand (for ad hoc riders) and commuter programs. However, current microtransit approaches face challenges in that resources (minivans and drivers) are expensive and, at the same time underutilized, while many ride requests are declined for lack of availability. Importantly, microtransit cannot ethically rely on price signals to manage high demand since the target user population is vulnerable and economically disadvantaged.

We consider a sociotechnical system (STS) (Singh, 2014) as a multistakeholder cyberphysical system. An STS has a social tier of people and organizations and a technical tier of cyberphysical resources and data. In the present setting, a municipal microtransit service constitutes an STS. Its stakeholders (including users and providers, i.e., riders, drivers, and the city transit authority) form the social tier of the STS. Its cyberphysical resources and data (i.e., vehicles and the associated information technology to request rides) form the technical tier of the STS. We posit that problems that may be difficult to solve at the technical tier can be made tractable through interventions in the social tier.

This paper describes an approach to improving the efficiency and effectiveness of microtransit by promoting prosocial attitudes among riders, reflected in their adjusting their preferences to facilitate sharing rides (Bardaka, Hajibabai, & Singh, 2020). We adopt a public-centric approach wherein we view riders as agents at the center of the multiagent microtransit system to ensure that the system is accepted by riders as safe, reliable, and trustworthy (Stein & Yazdanpanah, 2023). Specifically, we investigate a combination of social science and AI techniques to determine the best interventions to persuade microtransit riders to behave prosocially. We report on findings from a focus group with stakeholders, a cognitive model to serve as a surrogate for rider behavior, and an algorithm to learn both contextual elements and persuasive strategies to improve prosociality.

To build an agent that suggests acceptable interventions, our first challenge is to understand our riders. Suggestions must be tailored to individual riders for maximum effect. Some people are willing to adjust their departure or arrival times but not their pickup and dropoff locations (temporal versus spatial flexibility). For other riders, it may be the opposite case. Some people may be willing to walk further at certain times of the day than at others. Similarly, individuals may vary in their willingness to adjust their trips based on who stands to benefit. All these factors must be considered while suggesting alternatives. For example, on a sunny day, a certain rider may be willing to walk further for pickup than on a rainy day. Our agent should recognize that a prompt to walk further on a rainy day may only antagonize the rider and is unlikely to be accepted. However, the same rider may be willing to adjust their time, and this is what we aim to learn and suggest. Similarly, a female rider may have a low tolerance for walking on the street at night due to safety concerns. Our agent should not suggest that she compromise for someone else and walk further in the dark because that would undermine her well-being. If a rider needs to go somewhere urgently, a suggestion to postpone their ride will not be accepted. One person may be more willing to compromise if doing so would benefit a senior citizen, while another may empathize with neurodiverse individuals. If data about fellow riders could be captured and shared in such a

way that privacy constraints are adhered to, this information could be used to persuade riders to compromise for someone they perceive to be in need of help.

The foregoing motivation leads us to examine the following research questions.

RQ_{tolerance} Can we learn riders’ spatial tolerances to suggest optimal spatial adjustments?

RQ_{empathy} Can we learn riders’ empathetic tendencies to persuade them to adjust?

RQ_{profile} Could considering rider profile data lead to a better (nonnaive) starting point?

1.1 Approach and Contributions

To understand riders, we apply established social science methods to elicit their needs (requirements, risk attitudes, and values). The requirements refer to preferences regarding their trips and their flexibility. The risk attitudes refer to their views of risks such as being late or walking in the dark. The values refer to empathy and helping others in need. As our computational method, we apply machine learning to learn to persuade riders to relax their requests in light of their risk attitudes and values. For the purposes of evaluation, we capture rider needs and responsiveness to various persuasive messages.

We make four main contributions. Firstly, we support the applicability of a prosocial approach to microtransit through focus groups conducted with riders and drivers of a currently functioning city-wide microtransit system in Wilson, North Carolina. Secondly, we use ArcGIS (Booth, Mitchell, et al., 2001), a suite of online geographical information system software, to build a prototype mobile app for microtransit. Thirdly, we describe a reinforcement learning approach to determine the most effective interventions to persuade riders to behave prosocially. Finally, we use a cognitive architecture, Adaptive Control of Thought–Rational (ACT-R), to create a realistic model of rider decision-making, which we use as a surrogate for a microtransit rider.

1.2 Novelty

Previous research on AI for transportation and urban mobility has not tackled the challenges we address here. Some of it accommodates only predetermined rider preferences, such as driver competence and vehicle safety (Schleibaum & Müller, 2020; Yousaf, Li, Chen, Tang, & Dai, 2014). Other research focuses on economic incentives (Cipolina-Kun, Stein, Yazdanpanah, & Gerding, 2022). We consider users (i.e., riders) as central to the system and eliminate economic incentives in favor of persuasion toward prosociality. Some researchers have considered ways to enhance the attractiveness of alternative options to users, such as goal setting, personalized messaging, social comparison, and gamification (Anagnostopoulou, Bothos, Magoutas, Schrammel, & Mentzas, 2018; Kormos, Sussman, & Rosenberg, 2021). We build on their ideas, but we expand persuasive messaging to accommodate considerations of empathy. In addition, we consider an adaptive approach to persuade riders while learning their preferences, risk attitudes, and values—and respecting those risk attitudes and values in the persuasions attempted.

1.3 Organization

The rest of this paper is organized as follows. Section 2 describes the importance of microtransit in the rural setting. Section 3 describes the results of a focus group conducted with users of Wilson, North Carolina’s city-wide microtransit system. Section 4 demonstrates our idea of prosocial interventions in microtransit with the help of a prototype mobile app. Section 5 describes the Cooperative Adaptive Ride Sharing (CARS) agent, and our approach to learning riders’ spatial

tolerances and persuasive factors. Section 6 introduces the cognitive architecture ACT-R and explains how we use its Python implementation, `pyactr`, to simulate a microtransit rider by modeling rider decision-making for the task of accepting or rejecting walking suggestions. Section 7 describes experiments with a set of ACT-R riders who have different characteristics and preferences, and shows how our approach fares in learning acceptable interventions for them. Section 8 summarizes our findings and lays out some possible areas of future work in the field. Appendices A, B, C, and D contain detailed rider demographic data (collected through focus groups), reproducibility details for our experiments, additional experiment results, and a sensitivity analysis with different model hyperparameters, respectively.

Figure 1 illustrates our setting and the plan of the paper. As envisioned, a fielded solution would involve riders (humans) engaging with the CARS agent through a mobile app. The agent would model riders and attempt to persuade them to be prosocial (as described below) to improve overall system performance and rider satisfaction. In our experiments, we use a simplified version of the agent (considering a simplified environment) along with simulated riders, and address the research questions $RQ_{tolerance}$, $RQ_{empathy}$, and $RQ_{profile}$ introduced in Section 1. The riders are simulated based on the ACT-R cognitive architecture, including some parameters based on the Social Value Orientation (SVO) literature.

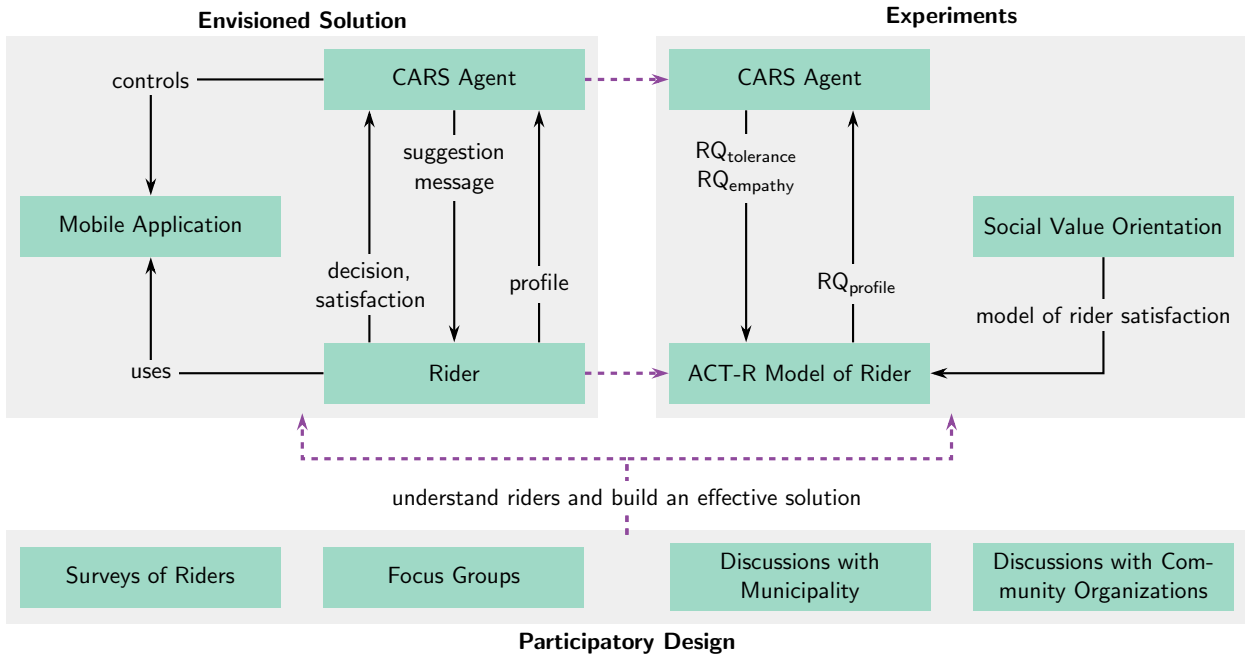


Figure 1: An illustration of our envisioned solution along with the research questions addressed in this paper as part of our experimental framework and our participatory design, intended to understand the needs of key stakeholders. Our experiments measure how far riders are generally willing to walk, how frequently they accept suggestions to walk to an alternative location, and their satisfaction with the adjustments they make.

2. Background on Microtransit

We consider the setting of (*public*) *microtransit* services that are emerging in rural areas in the US. Such services offer point-to-point on-demand rides, much like commercial services such as Uber or

Lyft. However, public microtransit often replaces fixed-route bus services and is often targeted at disadvantaged residents to enable them to access essential civic services. Thus, crucially, it does *not* rely on price signals.

Transportation is essential for residents to go to work, benefit from crucial services such as healthcare, and participate in civic life. People with disabilities or those who are elderly or poor must rely on public transportation. Rural areas have a low population density and fixed-route transit services (such as bus and rail) prove unviable since they are both expensive and underutilized. As a result, municipalities such as Wilson, North Carolina (our partner in this study) have shut down their fixed route transit and replaced it with microtransit through a small fleet of minivans, each able to hold a driver and up to six passengers.

Microtransit demand has been rising in small, disadvantaged communities due to its affordability, as microtransit rides are meant to be shared and only charge a nominal fee. Wilson, with a population of 40,000, was the first in North Carolina to implement a city-wide microtransit system, called RIDE, operated by Via (RIDE’s service provider). RIDE trips cost \$2.50 each, and riders can bring an extra person along for \$1.00. Children under the age of eight can travel for free, and senior citizens and disabled people are eligible for a discounted rate. Most riders book their trip through Via’s smartphone app, while reservations can also be made through a call center. The app provides convenient trip scheduling, trip updates, and digital fare payment.

Wilson is a community with a substantially higher proportion of households below the poverty line, zero-vehicle households, and people of color, than the North Carolina average. Small cities and towns like Wilson also have a strong sense of community and empathy between their members. Microtransit is the only public transit available in Wilson and receives about 20,000 trip requests per month. Most of the trips scheduled originate from the poorer neighborhoods of Wilson and go to the downtown area.

During a workshop we conducted with the key stakeholders of RIDE, we learned that unfortunately, during the morning and afternoon peak periods, a substantial fraction of ride requests received cannot be served. This is a major problem because, based on a survey conducted by Via, about 60% of the riders in Wilson use microtransit mainly for work and medical appointments. In addition, 86% are carless and 57% earn less than \$25K per year. Hence, many riders face daily struggles with the microtransit but cannot switch to other modes due to a lack of alternative travel options.

Despite the high demand, the microtransit vans (which can fit up to six passengers) remain highly underutilized. The system currently serves less than four trips per vehicle hour, and only one-third of the trips are shared with another booking. Additionally, RIDE does not support booking in advance.

3. Understanding Stakeholder Needs

We conducted interviews with all the key stakeholders. One group consists of the operational transportation managers in Wilson, from whom we learned about the economic constraints on the service.

A second group is of drivers in the current service. In an interview, the RIDE drivers described their experiences with riders. They stated that though some riders may be willing to walk, some, e.g., those with bags of heavy groceries, may not. Additionally, they stated that pickup and dropoff locations are not always convenient for riders, sometimes creating unsafe crossings and walk requests, as the current application does not consider local traffic and network conditions.

The largest and most important group of stakeholders in our setting is riders. We conducted focus groups in Wilson to understand this group. Unlike one-on-one interviews, focus groups

reveal the similarities and differences between participants in a social setting (Morgan, 2019). Semistructured focus groups are particularly useful for studies of how people make sense of a particular phenomenon or experience (Marshall & Rossman, 2021; Maxwell, 2012). 165 microtransit riders signed up for the five focus group sessions we organized (eIRB# 25553). We conducted five focus group sessions for a total of 32 participants, selected at random from the 165 candidates. Each participant received a research incentive of \$100 for a session of about one hour plus travel time. We invited six or seven participants to each session to ensure that they had space for free-flowing conversation.

Most participants arrived at our sessions by microtransit. The sessions took place at a local community organization’s conference center in downtown Wilson, where it was easily accessible to participants. In each session, participants completed a short survey, followed by 60 minutes of discussion. The sample was predominantly female (63%) and diverse (77% African American, 3% Caucasian, 3% American Indian/Alaskan Native, 7% Asian American/Pacific Islander, and 3% Other—Multiracial) with a mean age of 49 (age range: 27–70). The participants indicated that they use microtransit for commuting to work (68%), going to doctor’s appointments (87%), and running daily errands (74%). Others mentioned using microtransit to fulfill family responsibilities such as “pick[ing] up my kids from daycare” or visiting relatives in the nursing home. The majority (97%) reported using microtransit two to five times per week. More demographic data of the participants can be found in Table 7 in Appendix A.

Although the waiting times for the current system (RIDE) given by the participants range between 30 minutes to an hour, some try to book their rides at an earlier time to offset any delays: “I have to be at work at 5:00, so I try to set my ride at least about 3:50 because I know sometimes it might have a delay due to other pickups.” But, individuals who do not have phone access while at work (cashiers, factory employees) expressed that this is not an option. Others described cycles of canceled and rescheduled rides: “I might get a ride that says 30 minutes for the person, ok, I’m sitting there, wait, wait, wait. Get down to 12 minutes, then cancel that and put me over to another person. That person I wait for 20 minutes. Gets down to five minutes, they cancel that then give me another person.”

Many participants expressed flexibility in their travel schedules such that they were willing to identify vulnerable others whose rides should be prioritized over their own: “[T]ime is on my side, man, I got all the time.”; “[I]f a person in a wheelchair has a doctor’s appointment, then that’s a priority.” Willingness to walk varied from “a block or two” to “half a mile” and depended on individual attributes, weather, and safety: “I’m low vision”; “Ain’t a fan of the rain”; “The doctor has told me that they want me to get some exercise.”

Most were willing to share personal information through the app, and a few emphasized that sharing personal information should be optional. When asked about a rewards program to recognize volunteers, some participants deemed it a good idea, but many others said it is not needed: “kindness does not cost anything”; “I wouldn’t really care. I mean, if someone needs help, I’ll try my best to help them out. It’s not really a matter of getting something in return.”; “[W]e don’t need no incentives, just to help somebody out.”. Table 1 provides some comments made by focus group participants showing attitudes of prosociality and flexibility, as well as constraints that would be limiting factors for them.

4. Solution Concept and Illustration via a Mobile App

Figure 2 shows the proposed operation of the entire system, which we dub *Cooperative Adaptive Ride Sharing* or *CARS*.

Table 1: Attitudes of prosociality, constraints, and preferences in microtransit expressed by focus group participants (riders).

Verbatim Comment	Attitude
“Time is on my side, man, I got all the time”	Prosociality, Flexibility: time
“If a person in a wheelchair has a doctor’s appointment, then that’s a priority”	Prosociality, Other-interest, Empathy: riders in a wheelchair
“Someone who’s in our family shelter and has two little kids and a stroller and RIDE doesn’t provide car seats, so they have to carry the car seats with them. That’s a little hard to walk to any distance”	Prosociality, Empathy: riders with little kids
“Kindness does not cost anything”	Prosociality
“I wouldn’t really care. I mean, if someone needs help, I’ll try my best to help them out. It’s not really a matter of getting something in return”	Prosociality
“We don’t need no incentives, just to help somebody out”	Prosociality
“I’m low vision”	Constraint: vision
“I used to have to walk for four, sometimes five blocks, and for somebody who has a bad leg, that’s a lot”	Constraint: walking
“I’m sitting here waiting. I’m saying I have disability where I can’t stand for a long time. I was at Chiefs [store] and it was raining that day”	Constraint: standing, Preference: avoid rain
“My doctor does want me to walk at times, but not too much”	Constraint: walking
“The doctor, at times, has told me that they want me to get some exercise, so that’s about a mile for me to walk per day”	Persuasive factor: health
“Ain’t a fan of the rain”	Preference: avoid rain
“Because it’s cold”	Preference: avoid cold
“[Young daughter] has to walk from the corner to the house, and it’s dark”	Preference: avoid dark
“Going to be stranded” [if wheelchair loses charge]	Constraint: wheelchair
“Work or got doctor’s appointments, stuff like that where I need to be on time”	Constraint: urgency

This work focuses on the shaded region, developing the CARS agent to understand users (riders) and produce effective and persuasive suggestions for a rider, given the current conditions of the environment, fellow riders, and the agent’s knowledge of rider preferences.

We have built a prototype mobile app for microtransit to demonstrate our idea. We use ArcGIS, a collection of online geographic system software (Booth et al., 2001) to perform the geospatial computations required to calculate candidate alternative locations. We consider multiple riders who request trips on the app. Riders are clustered together based on the similarity of their routes, and an optimal route for sharing rides is computed. We then encourage riders to walk to a pickup point to avoid excessive detours. Riders may have a disability, in which case the algorithm will not suggest any alternative pickup point. Our prototype app includes support for the driver’s

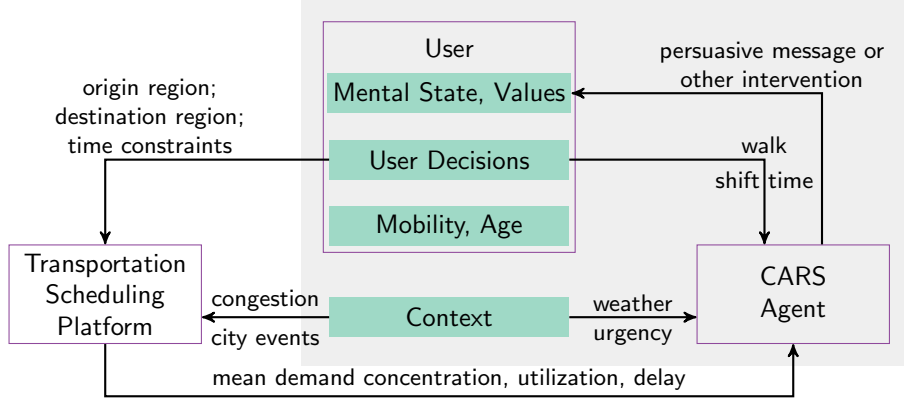


Figure 2: Proposed system operation with the focus of this study highlighted.

perspective but we omit it here since driver behavior is not relevant to this study. We note that all currently operating microtransit services allow for trip scheduling through a mobile app. However, none of the existing microtransit services enable or encourage users to show flexibility about trip pickup time or location. This app demonstration uniquely incorporates prosocial interventions in a microtransit scheduling platform.

4.1 Rider Pickup and Dropoff Locations

Riders can choose their pickup and dropoff points on the home screen after logging in, as shown in Figure 3. The locations can also be adjusted by moving the green or red pin, respectively. After choosing the pickup and dropoff locations, riders can request a ride by pressing the REQUEST RIDE button on the bottom right of the screen.


4.2 Suggesting an Alternative Location

We calculate the base route by considering the two farthest points in the cluster of pickup and dropoff locations and computing the route between them, considering ordered pairs of [pickup, dropoff] locations. For each rider, we compute the alternative location as the closest point on the base route from their requested location. Riders without disabilities will be suggested this alternative location, which they can accept or reject. The alternative pickup location is shown by a blue pin in Figure 4.


The rider acknowledges this alternative pickup location before proceeding. The walking path between the original pickup point (the green pin) and the alternative pickup point is shown in Figure 4 (b) as a black dotted line. Riders can choose to accept or reject the alternative location.

If the rider accepts the suggestion, the route is recalculated with the alternative pickup point, as shown in Figure 5 (a). The blue path depicts the final route to be taken by the driver. In future work, we imagine that a rider who accepts a suggestion would accrue KARMA POINTS. Those points could be used to gamify the app: to prioritize riders for timing and convenience (e.g., door-side pickup in times of need).

In case the suggestion is not accepted by the rider, the original path is used. As shown in Figure 5 (b), the route moves into the side road to pick up the rider. This would also happen in the case the rider has a disability, as in that case, the rider is picked up at their requested location.

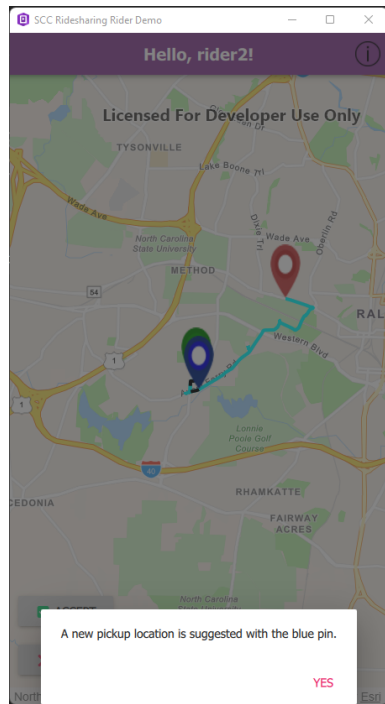


(a) Pickup

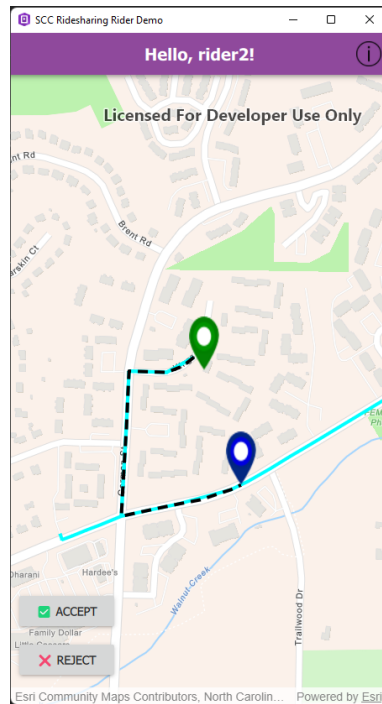


(b) Dropoff

Figure 3: Rider pickup and dropoff locations.



(a) Alternative Point Calculated



(b) Alternative Point Details Displayed

Figure 4: Presenting a suggestion to a rider.

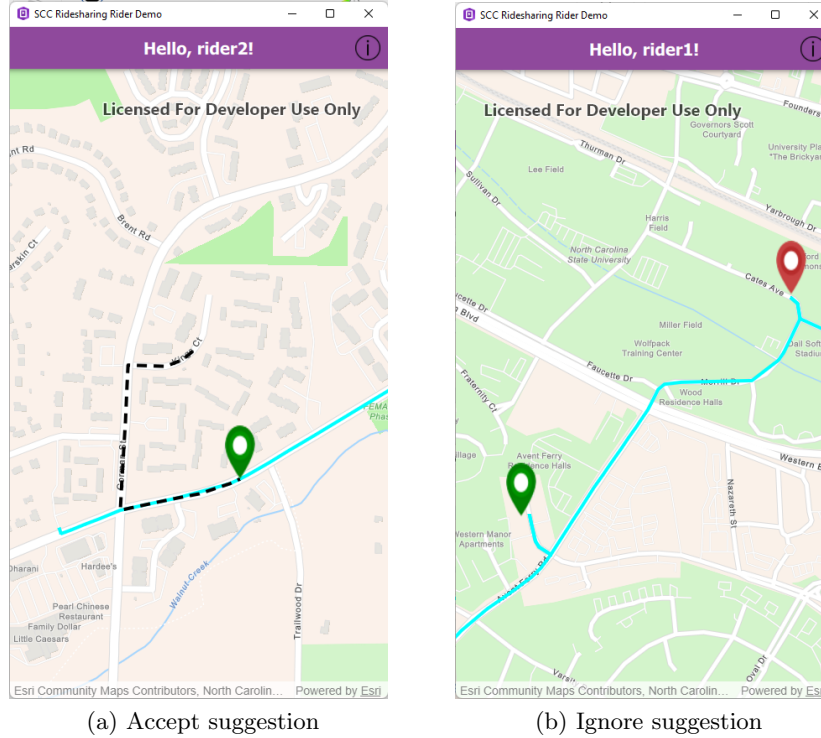


Figure 5: Rider responding to a suggestion.

5. CARS Agent

Rider preferences play a large role in the suggestions they accept. We aim to learn these preferences in two dimensions: a rider’s spatial (walking) tolerance under different environmental conditions (contexts), and the persuasive strategies that they respond to. For the purposes of this study, we consider a simplified environment with two features:

- Weather: sunny or rainy
- Time of day: morning, afternoon or evening

We presume that riders have certain spatial tolerances for each of the six environmental combinations produced from these features. We also assume riders have specific persuasive *value phrases*, corresponding to the categories of people or the general factors that the rider empathizes with. We describe a reinforcement learning approach to learn rider preferences in these two dimensions.

5.1 Spatial Adjustment Learning

We use model-free reinforcement learning (Proximal Policy Optimization) to learn optimal spatial suggestions for riders. We experiment with two models (with different reward functions) trained through interactions with the rider, and a customized model trained only on rider profile data.

We experiment with two reward functions, one that considers only the magnitude of accepted spatial adjustment by the rider, and one that considers rider satisfaction as well. Our aim is to learn how far a rider would be willing to walk under different environmental conditions (contexts).

5.1.1 CUSTOMIZED MODEL (PROFILE)

This model is trained on rider profile data, i.e., the preferences specified by a rider in their profile. For example, most people may be willing to walk further in the sun than in the rain. However, humans are notoriously incapable of *quantifying* their preferences. We suggest an alternative mode of initial input on the part of the rider: a *feature trace* (Bobu, Wiggert, Tomlin, & Dragan, 2022) specifying their weather preferences (for example, *sunny* > *rainy*), and their time of day preferences (for example, *morning*, *afternoon* > *evening*). We also consider other data provided by the rider (their age and gender). Our aim is to start out with some understanding of the rider with minimal input from them. For example, the model should know not to suggest longer distances to walk when it is raining or to walk in the dark if the rider has already indicated they would rather not. Using this rider profile, we create a default ordering of environmental conditions and set base expected spatial tolerances for different contexts accordingly.

An estimation of the rider’s spatial tolerance can be calculated as a function of their profile data as shown in Equation 1, where c is the context (environmental state), comprised of weather and time of day.

$$\text{estimated-spatial-tolerance}[c] = F(\text{age}, \text{gender}, \text{feature-trace}[c]) \quad (1)$$

The customized model is trained to learn these estimated spatial tolerances. It suggests spatial adjustments and is rewarded based on the magnitude of adjustment suggested, given that this magnitude is not more than the estimated tolerance in the context. In the following equation, s and c represent spatial adjustment suggested and context, respectively.

$$\text{reward}[s, c] = \begin{cases} s, & \text{if } s \leq \text{estimated-spatial-tolerance}[c] \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

5.1.2 SPATIAL MODEL (SPATIAL)

This model is trained through interactions with the rider, with a reward function equal to the magnitude of accepted spatial adjustment by the rider. In the following equation, s and c are the spatial adjustment and context as in the customized model.

$$\text{reward}[s, c] = \begin{cases} s, & \text{if } s \text{ is accepted by the rider} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

5.1.3 RIDER SATISFACTION MODEL (SAT+SPATIAL)

This model is trained through interactions with the rider, with a reward function equal to a combination of both the magnitude of accepted spatial adjustment and rider satisfaction.

This model assumes that riders can provide an accurate measure of their satisfaction when they accept a suggested adjustment. In the following equation, c , s , and r are the context, spatial adjustment suggested, and rider satisfaction with the spatial adjustment s respectively.

$$\text{reward}[s, c] = \begin{cases} s \times r, & \text{if } s \text{ is accepted by the rider} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

5.2 Persuasive Strategy Learning

The second preference we wish to learn is *who* and *what* riders are willing to adjust for. For the purposes of this study, we consider three persuasive components for the implementation of a persuasive strategy: empathy for those in need, environmental benefit, and health benefit to oneself. We consider six categories of people that could be presumed to be in need of support: *babies, children, senior citizens, ill people, people with disabilities, and neurodiverse individuals*. We presume that riders have preconceived notions regarding the people they empathize with. People may also empathize with *environmental benefit* and *health benefit* (due to saving fuel costs and walking, respectively). We refer to these categories of people and factors as *value phrases*. We assume that each rider may be empathetic toward a subset of these value phrases.

In our proposed system, each of these phrases will be mapped to one or more sentences. For example, *environmental benefit* could be mapped to “You will be saving x gallons of fuel!” or “You will be helping to prevent the emission of y units of carbon dioxide!”, *health benefit* could be mapped to “By walking x meters to location A, you would be getting y steps, the first step to healthy living!”, and *senior citizens* could be mapped to “Would you be willing to walk x meters further to location B? You would be helping a senior citizen with their trip.”

In our experiments, we simulate current conditions of the environment during each ride, i.e., the fellow riders and general factors. If the rider rejects the spatial adjustment, we attempt to persuade them further by suggesting who or what they would be helping by making this compromise. We want to learn what is most persuasive to the rider to maximize the chance of them accepting our suggestion. We use a multiarmed bandit algorithm to learn the rider’s empathetic preferences based on their response to this second prompt. Over time we learn the value phrases each rider is empathetic toward, and so we know what to say to the rider to maximize our chances of them accepting the suggestion.

6. Cognitive Modeling

To build a realistic model of rider decision-making to represent riders in our experiments, we use a cognitive architecture, i.e., a psychological model of human cognition upon which specific tasks can be defined (Anderson & Lebiere, 2003). The underlying psychological components are shared between tasks. We adopt the cognitive architecture ACT-R (Adaptive Control of Thought—Rational) (Anderson, 1996) because it has an easy-to-use Python library (Brasoveanu & Dotlačil, 2020). Accordingly, the riders in our experiments that the CARS agent interacts with are defined using ACT-R. Currently, we do not model drivers in our experiments.

6.1 Overview of ACT-R

In ACT-R (Anderson, 1996), the human mind is modeled as a set of three components or cognitive modules, which work together to process information and produce behavior. These three components are the *declarative memory*, which stores facts; the *procedural memory*, which defines actions; and the *control system*, which coordinates the interaction between the declarative and procedural memory and generates behavior according to environmental state. Figure 6 describes ACT-R schematically. To reduce clutter, we removed the information flows, which go into and out of the Productions box. The visual buffer processes visual information from the world and the manual buffer makes changes in the world. Data in ACT-R is represented in the form of chunks, mimicking the way humans process, store and retrieve small pieces of information.

Decision-making in ACT-R works by way of utilities estimated for different productions. Procedural memory defines the available productions and their prerequisites and consequent states.

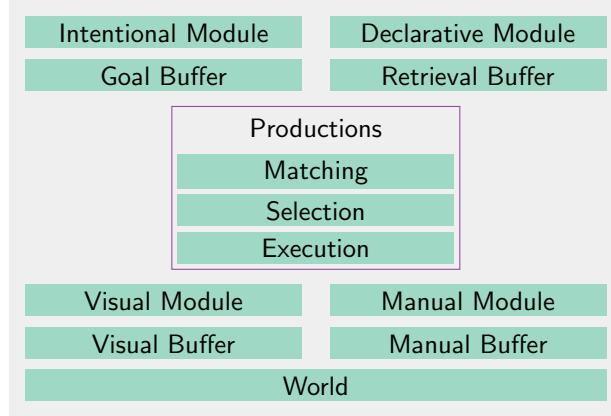


Figure 6: Modular software architecture of the cognitive architecture ACT-R.

Given the world state (represented by certain chunks in the goal buffer), we define productions (actions) that can be carried out.

After a production is fired, ACT-R carries out the transitions until the final state, and the utilities of the productions leading up to that state are updated based on the final reward. Over time, the productions which lead to a reward have a higher utility, and hence the ACT-R rider ‘prefers’ those productions given a choice of multiple productions for a certain set of chunks in the goal buffer.

Utility updates in ACT-R occur according to the following formula:

$$U(i, t + 1) = U(i, t) + \alpha[R(t) - U(i, t)]$$

where:

$U(i, t)$ is the current utility of production rule i at time t .

$U(i, t + 1)$ is the updated utility of production rule i at time $t + 1$.

α is the learning rate, which determines the extent to which the utility is updated in response to feedback from the environment.

$R(t)$ is the reinforcement signal or reward received by the system at time t , which reflects the success or failure of the action taken based on the production rule.

$[R(t) - U(i, t)]$ is the prediction error, which represents the difference between the utility of the production rule and the actual reward received by the system.

6.2 Social Value Orientation

Social value orientation (SVO) is a concept in social psychology that states that different individuals have different preferences regarding the allocation of resources between themselves and others (Lange, 1999). Social value orientation affects behavior in social dilemmas (Balliet, Parks, & Joireman, 2009). For the task of adjusting for the sake of others with respect to microtransit rides, social value orientation hence plays a large role, so we use SVO as a part of the internal rewards in the ACT-R riders.

A person’s social value orientation is reflected in the internal value they acquire from performing actions. For example, a competitive person would prioritize their outcomes and prefer others’ outcomes to be worse than theirs. An altruistic person would care more about the outcomes of others. Some people may prioritize themselves but also prefer that others get what they want.

Correspondingly, the internal satisfaction a rider gets from an action (accepting or rejecting an adjustment that benefits someone else) changes based on their social value orientation.

We adopt a simplified version of the classic SVO framework (Murphy & Ackermann, 2014), modeling an ACT-R rider’s SVO using two parameters.

Other-interest is the degree to which the rider values the outcome of others relative to their own. It ranges from 0 (no interest in the outcome of others) to 1 (only interested in the outcome of others). The complement of other-interest is *self-interest* (i.e., self-interest = $1 - \text{other-interest}$).

Prosociality is the degree to which the rider values the sum of outcomes for themselves and others. It ranges from 0 (completely competitive) to 1 (completely prosocial). The complement of prosociality is *competitiveness* (i.e., competitiveness = $1 - \text{prosociality}$).

6.3 Persuasive Value Phrases

As mentioned in Section 5.2, we consider three components that could persuade riders to accept a previously-rejected suggestion. These are empathy for those in need, environmental benefit, and health benefit to oneself.

We refer to these as *value phrases* and assume that each rider may be empathetic toward a subset of these value phrases. Even proself riders can be persuaded to adjust if we know what they are empathetic toward.

To preserve realism, we encode riders to have responses based on similarity to the causes they believe in. For example, if they sympathize with children, they would be more likely to be empathetic toward babies as well. For this, we use a language model to calculate the cosine similarity between the value phrase suggested and the value phrases the rider responds to.

6.4 Spatial Tolerances

For the purposes of this study, we model riders to have certain spatial tolerances for each of the possible contexts ([weather, time of day] combinations) mentioned in Section 5. A tolerance is not a threshold, but does affect the internal rewards riders acquire from different productions. A person with a tolerance of 100 meters in a certain context would respond differently to an adjustment of 120 meters compared to someone with a tolerance of 50 meters, since the amounts of adjustment beyond their tolerances they would have to make are different. We use this assumption to define riders’ internal rewards for productions, combined with other factors.

6.5 Rider Internal Reward Definition

For a realistic utility update of riders’ actions, we must define rewards that accurately reflect the internal value a rider gets from a certain action, i.e., the satisfaction they acquire. We model this process as goal-directed choice or value-based decision-making (Rangel, 2009). Riders make decisions based on a comparison of utilities, reflecting the satisfaction they are likely to receive from their choices.

For the purposes of this study, we consider three cases of a spatial adjustment suggestion, corresponding to a low (less than 25% above their spatial tolerance), medium (between 25% and 75% above their spatial tolerance), and high (more than 75% above their spatial tolerance) amount of adjustment on the part of the rider. We assume that the rider’s internal reward achieved by accepting these suggestions is inversely proportional to the amount of adjustment. The internal reward attained by rejecting suggestions is directly proportional to the amount of adjustment:

riders are likely to feel worse about rejecting interventions that inconvenience them less. In other words, we assume that:

$$\begin{aligned} \text{rider-internal-reward}[\text{low-adjustment-accept}] &> \text{rider-internal-reward}[\text{medium-adjustment-accept}] \\ &> \text{rider-internal-reward}[\text{high-adjustment-accept}] \\ \text{rider-internal-reward}[\text{low-adjustment-reject}] &< \text{rider-internal-reward}[\text{medium-adjustment-reject}] \\ &< \text{rider-internal-reward}[\text{high-adjustment-reject}] \end{aligned}$$

We calculate rider-internal-reward with the following formula.

$$\text{rider-internal-reward} = \begin{cases} 1/a \times o \times p \times m \times (1 - u) & \text{if accept} \\ a \times (1 - o) \times (1 - p) \times (1 - m) \times u & \text{if reject} \end{cases} \quad (5)$$

where

- a indicates the *adjustment*, i.e., the fraction that the rider would need to walk extra above their tolerance.
- o indicates the *other-interest*, i.e., the degree to which the rider values the outcome of others relative to their own; o ranges from 0 (no interest in the outcome of others) to 1 (only interested in the outcome of others).
- p indicates the *prosociality* of the rider, indicating the degree to which the rider values the sum of outcomes for themselves and others; p ranges from 0 (completely competitive) to 1 (completely prosocial).
- m indicates the *mood* of the rider, affecting their response to a certain suggestion at a certain time. This adds an additional factor of randomness in the response, mimicking the unpredictable nature of human behavior. It ranges from 0 (completely negative) to 1 (completely positive).
- u indicates the *urgency* of the current trip. It ranges from 0 (not urgent, implying more flexibility on behalf of the rider) to 1 (extremely urgent).

7. Experiments and Results

We aim to learn two types of preferences for riders: their spatial tolerances in different contexts and the factors they empathize with. Our CARS agent combines this knowledge to persuade riders to accept spatial adjustments and behave prosocially. To assess the performance of our approach in rider modeling, we run experiments with five diverse riders. A rider’s characteristics (their SVO orientation) and their current state while requesting a ride (their mood and the urgency of the trip) affect their response to suggested adjustments in addition to their actual spatial tolerances and persuasive value phrases.

7.1 Rider Personas

We consider four archetypal personalities (competitive, individualistic, prosocial, and altruistic) according to SVO research as well as one moderate persona to offset the extremities. Competitive

Table 2: Rider personas for this study.

SVO	Gender	Age	Other-interest	Prosociality
Competitive	Male	Early 30s	0.01	0.01
Individualistic	Female	16	0.01	0.50
Moderate	Female	Mid 40s	0.45	0.50
Prosocial	Male	Early 50s	0.50	0.99
Altruistic	Female	Early 70s	0.99	0.99

people seek to maximize the difference between their outcomes and those of others, i.e., to maximize their outcomes while minimizing others’. Individualistic people are concerned only with maximizing their own outcomes. Prosocial people prefer mutually beneficial outcomes, considering both their own outcomes and the outcomes of others. Altruistic people are interested only in the outcomes of others—they are willing to sacrifice their own well-being for the benefit of others. In addition to these four archetypal SVOs, we consider one moderate persona who prioritizes their outcomes slightly more than others’, and is neither prosocial nor competitive. We consider that prosociality tends to increase with age (Matsumoto, Yamagishi, Li, & Kiyonari, 2016) while modeling these personas. We also model riders above the age of 65 to have slightly lower spatial tolerances for all times of the day, and females to have lower spatial tolerances in the evening. A summary of our rider personas is shown in Table 2. We show sample rider profile data and internal preferences (those used for the prosocial rider in our experiments) in Tables 3 through 5. The remaining riders’ details can be found in Appendix B.2.

7.2 Hypotheses

The following hypotheses are evaluable claims based on our research questions.

H_{tolerance} *Spatial tolerances for riders can be learned over time.* This knowledge can be used to suggest optimal spatial adjustments that would benefit the system without causing too much inconvenience to riders.

H_{empathy} *Persuasive strategies to promote prosociality can be learned over time.* This can be used to persuade riders to behave prosocially even if the spatial adjustment causes them some inconvenience.

H_{profile} *Knowledge about a rider’s basic profile will provide a better (nonnaive) starting point to interact with the rider, i.e., we can use basic rider profile data to initialize a nonnaive (customized) model to start with.*

7.3 Evaluation Metrics

We evaluate the performance of the CARS agent as a combination of the performance of the spatial tolerance learning and the persuasive strategy learning aspects of it.

7.3.1 CARS AGENT–SPATIAL TOLERANCE LEARNING

We use the *average accepted spatial adjustment per episode* and the *average acceptance percentage per episode* as our evaluation metrics. Episodes are of a fixed length of 1,000 time steps. We evaluate

Table 3: Sample rider profile: Prosocial rider

Parameter	Value
Gender	Male
Age	Early 50s
Weather feature trace	Sunny > Rainy
Time of day feature trace	[Morning = Evening] > Afternoon

Table 4: Sample spatial tolerances: Prosocial rider

Weather	Time of Day	Spatial Tolerance (in meters)
Sunny	Morning	475
Sunny	Afternoon	200
Sunny	Evening	400
Rainy	Morning	250
Rainy	Afternoon	200
Rainy	Evening	20

the three models mentioned in Section 5.1. PROFILE, on interacting with the rider, converges to the same policy as SPATIAL. The longer the customization, the longer the model takes to converge to this policy. We hence evaluate the performance of PROFILE (trained for 50,000 time steps), SPATIAL (trained for 150,000 time steps), and SAT+SPATIAL (trained for 150,000 time steps).

7.3.2 CARS AGENT-PERSUASIVE STRATEGY LEARNING

We evaluate the performance of our agent by calculating the similarity between the actual and predicted value phrase distributions for the rider.

7.4 Results

We show results for spatial adjustment learning for all the riders in Figures 7, 8, 9, and 10. We experimented with both the default PPO hyperparameters (Table 10) and modified hyperparameters (varying the learning rate and n_steps). We used the default hyperparameters for SPATIAL and PROFILE, and a learning rate of 0.003 for SAT+SPATIAL. This is because SAT+SPATIAL showed

Table 5: Sample empathy (persuasive value phrase) distribution: Prosocial rider

Value Phrase	Persuasive Percentage
Environmental benefit	51%
People with disabilities	49%




Figure 7: Average Accepted Adjustment for the three models and the baseline (untrained), measured over 500 episodes of 1,000 time steps each. SAT+SPATIAL outperforms the other models, getting riders to walk further, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.

a marked improvement with these hyperparameters, although the other models did not. Rider satisfaction was also considered in selecting the best hyperparameters for each model. Table 25 in Appendix D provides detailed results for all models we experimented with, and also shows total rider satisfaction for each model.

We show the difference between actual and predicted persuasive value phrase distributions in Table 6, and Figure 11 shows the actual versus top three predicted value phrase percentages for the prosocial rider. More results for persuasive value phrase learning can be found in Appendix C.2.

7.4.1 EVALUATING HYPOTHESIS $H_{\text{TOLERANCE}}$

Figures 7 and 8 show spatial adjustment results in terms of the average accepted adjustment and acceptance percentage for each rider for the three models and the untrained baseline, measured over 500 episodes of 1,000 time steps each. There is a slight improvement in PROFILE compared to the untrained model, indicating the benefit of customization with the rider profile. SPATIAL, which has been trained for 150,000 time steps with the rider, performs only slightly better than PROFILE. However, SAT+SPATIAL, which has been trained for the same number of steps with the rider, performs better than the other models in terms of both average accepted adjustment and average acceptance percentage. This indicates that considering rider satisfaction is beneficial for the system as well, as riders can be persuaded to make larger adjustments if we consider their satisfaction while accepting adjustments. The improvement in SAT+SPATIAL is much more pronounced for the Competitive, Individualistic, and Moderate riders, indicating that even riders who are proself can behave prosocially if we understand them better. Most people are not completely prosocial (as are the Prosocial and Altruistic personas, who, being completely prosocial, accept almost all suggestions even if they are highly inconvenienced by them), and so are likely to be much more satisfied (and willing to adjust) if we take their satisfaction into account.

Figures 9 and 10 show the Cohen’s d effect sizes and corresponding 95% confidence intervals of SAT+SPATIAL compared to SPATIAL, PROFILE, and the untrained baseline. More detailed results can be found in Appendix C.1.




Figure 8: Average Acceptance Percentage for the three models and the baseline (untrained), measured over 500 episodes of 1,000 time steps each. SAT+SPATIAL outperforms the other models, getting riders to behave prosocially more frequently, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.




Figure 9: Effect sizes measured as Cohen's d scores and corresponding 95% confidence intervals of the improvement of SAT+SPATIAL over SPATIAL, PROFILE and an untrained model for accepted adjustment value. This is measured over 500 episodes, each of length 1,000 time steps. SAT+SPATIAL outperforms the other models, getting riders to walk further, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.

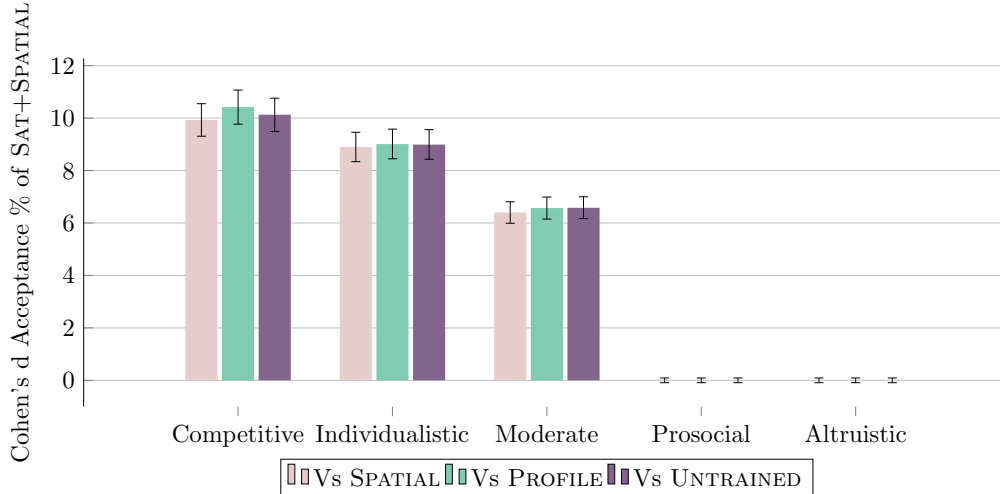


Figure 10: Effect sizes measured as Cohen’s d scores and corresponding 95% confidence intervals of the improvement of SAT+SPATIAL over SPATIAL, PROFILE and an untrained model for the percentage of accepted suggestions. This is measured over 500 episodes, each of length 1,000 time steps. SAT+SPATIAL outperforms the other models, getting riders to behave prosocially more frequently, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.

7.4.2 EVALUATING HYPOTHESIS H_{EMPATHY}

Figure 11 provides a summary of the persuasive value phrase distribution learned by our bandit model for the prosocial rider, and Table 6 shows the Hellinger distance between the actual and predicted persuasive value phrase distributions for each of the experiments. The Hellinger distance quantifies the difference between two probability distributions and is defined for two discrete probability distributions $P = (p_1, \dots, p_k)$ and $Q = (q_1, \dots, q_k)$ as

$$H(P, Q) = \frac{1}{\sqrt{2}} \times \sqrt{\sum_{i=1}^k (\sqrt{p_i} - \sqrt{q_i})^2}. \quad (6)$$

We use the Hellinger distance to calculate distribution similarity because it is a symmetric and bounded measure of the difference between two probability distributions. The Hellinger distance can range from 0 (perfect similarity) to 1 (maximum dissimilarity).

Due to noise during learning (other conditions such as spatial tolerance for the rider given the environmental conditions, and randomness mimicking human behavior implemented as a part of ACT-R), the results are not exact, but on the whole, our model is able to successfully learn probabilities of value phrases being persuasive to riders. The persuasive value phrase distributions for the other riders can be found in Appendix C.2.

7.4.3 EVALUATING HYPOTHESIS H_{PROFILE}

With the current rider profile data, there is a slight improvement in PROFILE compared to the untrained baseline, as shown in Figures 7 and 8 and in Table 23 in Appendix C.1.




Figure 11: Prosocial rider: actual vs top three predicted persuasive value phrase percentages; the remaining 11% is divided between the rest of the phrases.

Table 6: Hellinger distance between actual and predicted persuasive value phrase distributions.

Experiment	Hellinger Distance
Prosocial	0.391
Altruistic	0.338
Moderate	0.604
Competitive	0.316
Individualistic	0.261

7.4.4 SUMMARIZED RESULTS

We now return to the research questions introduced in Section 1 and summarize how we answer them in this paper.

RQ_{tolerance} *Can we learn riders’ spatial tolerances to suggest optimal spatial adjustments?*

We answer this question positively based on Figures 7, 8, 9, and 10 in Section 7.4.1.

RQ_{empathy} *Can we learn riders’ empathetic tendencies to persuade them to adjust?*

We answer this question positively based on Table 6 in Section 7.4.2 and Figure 11 in Section 7.4.2. Figures 12 through 15 in Appendix C.2 give additional details.

RQ_{profile} *Could considering rider profile data lead to a better (nonnaive) starting point?*

We answer this question positively based on Figures 7 and 8 in Section 7.4.1. Table 23 in Appendix C.1 gives additional details.

8. Discussion

We present a conception of a prosocial approach to microtransit to create a more equitable and sustainable ecosystem. We demonstrate the working of our idea with a prototype app, as well as experiments using a cognitive architecture as a surrogate for a human rider, to show that rider preferences can be learned with reinforcement learning and used to persuade them to help others. We show that both acceptable spatial adjustments and persuasive strategies can be learned. We show that minimal input from the rider (of their basic profile data) is enough to customize a model with some knowledge to start with, so our rider is faced with less inconvenience at the time of their first ride request.

We find that the methods are effective for riders with varying levels of prosociality. Fortunately, as we found through the focus groups, many riders are prosocial at the outset. Considering rider satisfaction can help persuade even the less prosocial riders to behave prosocially.

We show that with ACT-R, we are able to model diverse riders, whose varied responses to suggestions indicate the differing internal satisfaction they derive from their actions. Our results show that ACT-R can be used to model human decision-making in simulations where human input is required and relevant or sufficient data is not available, accounting for people with different behaviors and motivations.

Our results suggest that if we combine this work with an optimizer that optimizes a system-level (for example, average trip time) or higher priority (for example, a more in-need rider’s convenience) metric, we could calculate adjustments that are of the least inconvenience to riders while increasing the prosociality of the entire system.

8.1 Challenges

Using AI-based interventions to change user preferences or behavior, even for a good societal objective, is potentially ethically risky. Key challenges include a deeper understanding of consent (Singh, 2022) and privacy requirements (Ogunniye & Kokciyan, 2023) so that an AI agent does not violate a user’s autonomy.

A more general challenge is that of achieving trust. A decision about trust brings forth judgments of an agent’s ability, benevolence, and integrity (Mayer, Davis, & Schoorman, 1995). The same constructs form an effective basis for assessing the trustworthiness of AI agents (Singh & Singh, 2023). Achieving trustworthiness in the present scenario would involve care for the following

kinds of concerns. For ability: the agent (and associated hardware such as vehicles) is competent and reliable, e.g., in finding safe pickup and dropoff points for riders. For benevolence: the agent acts in a user’s interest, e.g., by informing them of all relevant options, including the best ones, so the user can choose freely. For integrity: the agent acts honorably, e.g., by protecting user privacy and not misusing user data. For trust, we would like not only the agent to be trustworthy but also the users to place the requisite trust in it. Besides the inherent benefits of ensuring that our STSs promote trust and that our agents are trustworthy, another motivation for trust is practical: Once a user loses trust in the system, they may elect not to participate or participate only to the extent necessary, e.g., by disregarding any attempted persuasion and thereby forgoing the prosocial outcomes we desire.

8.2 Future Work

This research opens up interesting directions for further study. First, we may explore the effects of more shared rider data with the agent. For example, we may explore the possibility of getting an initial approximation of riders’ SVOs when they sign up for the application, which would provide our model with an understanding of the rider at the very start and lead to better performance of the customized model. We could also explore a multiobjective reinforcement learning approach for the problem, considering system benefit, spatial preferences for riders, and persuasive strategies for riders as separate objectives to be optimized. We may also explore the combination of the preference learning shown in this work with an optimizer to compute the resulting effect on system conditions and attest to whether it contributes to overall prosociality.

Second, the same idea can be applied to persuading riders to adjust their requested times. We can learn rider preferences for adjusting their times in a comparable way, with the only difference being that the urgency of the trip is likely to be more important in temporal flexibility. Given knowledge of a rider’s spatial and temporal flexibility under certain conditions, we can suggest the optimal intervention: one that they are likely to accept and that is most beneficial to the system as a whole. If we use an optimizer algorithm for a certain cluster of riders and drivers, considering real-time traffic conditions, weather, and other environmental factors, these preferences will help decide which interventions are optimal for the given ride.

Our solution includes ideas such as allowing riders to book in advance. This may lower cancellation rates, as riders may be less likely to cancel when their waiting times are reduced. Such riders would prefer to book in advance for a trip that they cannot be flexible with (such as going to work or a healthcare appointment) as well as a trip they can be flexible with (such as going to a grocery store). Our current project (Bardaka, Hentenryck, Lee, Mayhorn, Monast, Samaranayake, & Singh, 2024) is investigating such ideas.

Rider empathy may also change based on environmental factors. For example, if it is heavily raining, people may be more empathetic toward others than they would be otherwise. If it is late, a male rider may be understanding of a female’s needs and be willing to compromise. Future work can enhance the complexity of such persuasive methods.

In addition to learning preferences and preconceptions of riders, other methods could be used to persuade them. Gamifying the system is one way to do so (Ro, Brauer, Kuntz, Shukla, & Bensch, 2017) where people can compete for karma points or play in teams for normative influence, which would be especially effective for those who tend to conform to society. We could also make use of choice architectures (Münscher, Vetter, & Scheuerle, 2016). Learning preferences is at the core of effective persuasion, however, once those preferences are learned, there are multiple ways to increase the attractiveness of alternative options to a rider. Future work can also explore these areas.

Acknowledgments

We thank Aditya Khadse for his effort in implementing our prototype. We thank the anonymous reviewers for their helpful comments. We acknowledge support from the US National Science Foundation (grants SCC-2125447 and SCC-2325720).

References

- Anagnostopoulou, E., Bothos, E., Magoutas, B., Schrammel, J., & Mentzas, G. (2018). Persuasive technologies for sustainable mobility: State of the art and emerging trends. *Sustainability*, *10*(7), 2128.
- Anderson, J. R., & Lebiere, C. (2003). The Newell test for a theory of cognition. *Behavioral and Brain Sciences*, *26*(5), 587–601.
- Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist*, *51*(4), 355–365.
- Balliet, D., Parks, C., & Joireman, J. (2009). Social value orientation and cooperation in social dilemmas: A meta-analysis. *Group Processes & Intergroup Relations*, *12*(4), 533–547.
- Bardaka, E., Hajibabai, L., & Singh, M. P. (2020). Reimagining ride sharing: Efficient, equitable, sustainable public microtransit. *IEEE Internet Computing (IC)*, *24*(5), 38–44.
- Bardaka, E., Hentenryck, P. V., Lee, C. C., Mayhorn, C. B., Monast, K., Samaranayake, S., & Singh, M. P. (2024). Empathy and AI: Achieving equitable microtransit for underserved communities. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), Special Track on AI for Good*, pp. 7179–7187, Jeju, Korea. IJCAI.
- Bobu, A., Wiggert, M., Tomlin, C. J., & Dragan, A. D. (2022). Inducing structure in reward learning by learning features. *International Journal of Robotics Research*, *41*(5), 497–518.
- Booth, B., Mitchell, A., et al. (2001). *Getting Started with ArcGIS*. ESRI New York.
- Brasoveanu, A., & Dotlačil, J. (2020). The ACT-R cognitive architecture and its pyactr implementation. In *Computational Cognitive Modeling and Linguistic Theory*, chap. 2, pp. 7–37. Springer International Publishing.
- Cipolina-Kun, L., Stein, S., Yazdanpanah, V., & Gerding, E. (2022). Coalition formation in ridesharing with walking options. In *ICLR 2022 Workshop on Gamification and Multiagent Solutions*. Accessed 2023-06-28.
- Kormos, C., Sussman, R., & Rosenberg, B. (2021). How cities can apply behavioral science to promote public transportation use. *Behavioral Science & Policy*, *7*(1), 95–115.
- Lange, P. A. M. V. (1999). The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation. *Journal of Personality and Social Psychology*, *77*(2), 337.
- Marshall, C., & Rossman, G. B. (2021). *Designing Qualitative Research*. SAGE Publications, Inc. Accessed 2023-06-28.
- Matsumoto, Y., Yamagishi, T., Li, Y., & Kiyonari, T. (2016). Prosocial behavior increases with age across five economic games. *PloS ONE*, *11*(7), e0158671.
- Maxwell, J. A. (2012). *Qualitative Research Design: An Interactive Approach*. Applied Social Research Methods. SAGE Publications, Inc. Accessed 2023-06-28.

- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *The Academy of Management Review*, *20*(3), 709–734.
- Morgan, D. L. (2019). *Basic and Advanced Focus Groups*. SAGE Publications, Inc.
- Murphy, R. O., & Ackermann, K. A. (2014). Social value orientation: Theoretical and measurement issues in the study of social preferences. *Personality and Social Psychology Review*, *18*(1), 13–41.
- Münscher, R., Vetter, M., & Scheuerle, T. (2016). A review and taxonomy of choice architecture techniques. *Journal of Behavioral Decision Making*, *29*(5), 511–524.
- Ogunniye, G., & Kokciyan, N. (2023). A survey on understanding and representing privacy requirements in the internet-of-things. *Journal of Artificial Intelligence Research*, *76*, 163–192.
- Rangel, A. (2009). The computation and comparison of value in goal-directed choice. In Glimcher, P. W., Camerer, C. F., Fehr, E., & Poldrack, R. A. (Eds.), *Neuroeconomics*, chap. 28, pp. 425–440. Academic Press, London.
- Ro, M., Brauer, M., Kuntz, K., Shukla, R., & Bensch, I. (2017). Making cool choices for sustainability: Testing the effectiveness of a game-based approach to promoting pro-environmental behavior. *Journal of Environmental Psychology*, *53*, 20–30.
- Schleibaum, S., & Müller, J. P. (2020). Human-centric ridesharing on large scale by explaining ai-generated assignments. In *Proceedings of the 6th EAI International Conference on Smart Objects and Technologies for Social Good*, pp. 222—225. Association for Computing Machinery.
- Shaheen, S., Cohen, A., Chan, N., & Bansal, A. (2020). Sharing strategies: Carsharing, shared micromobility (bikesharing and scooter sharing), transportation network companies, micro-transit, and other innovative mobility modes. *Transportation Land Use and Environmental Planning, Berlin, Germany: Elsevier*, *1*, 237–262.
- Singh, A. M., & Singh, M. P. (2023). Wasabi: A conceptual model for trustworthy artificial intelligence. *IEEE Computer*, *56*(2), 20–28.
- Singh, M. P. (2014). Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology*, *5*(1), 1–23.
- Singh, M. P. (2022). Consent as a foundation for responsible autonomy. *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, *36*(11), 12301–12306. Blue Sky Track.
- Stein, S., & Yazdanpanah, V. (2023). Citizen-centric multiagent systems. In Agmon, N., An, B., Ricci, A., & Yeoh, W. (Eds.), *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pp. 1802–1807. Association for Computing Machinery.
- Yousaf, J., Li, J., Chen, L., Tang, J., & Dai, X. (2014). Generalized multipath planning model for ride-sharing systems. *Frontiers of Computer Science*, *8*, 100–118.

Appendix A. Focus Group Demographic Data

Table 7 provides detailed rider demographic data, captured during the focus group we conducted.

Table 7: Focus Group Demographic Data (Riders).

Variable	Categories	N	%
Gender	Male	11	36.7
	Female	19	63.3
Race	African American	24	77.4
	Caucasian	1	3.2
	American Indian/Alaskan Native	1	3.2
	Asian American/Pacific Islander	2	6.5
	Other - Multiracial	1	3.2
Education	Less than high school	3	10.0
	High school graduate	7	23.3
	Some college	14	46.7
	Bachelor's degree	5	16.7
	Master's degree	1	3.3
Housing	House/Apartment/Condo	27	90.0
	Senior housing	1	3.3
	Relative's home	2	6.7
Self-health rating	Fair	9	29.0
	Good	11	35.5
	Very good	7	22.6
	Excellent	4	12.9
Physical disability	Yes	14	45.2
	No	17	54.8
Mobile impairment	Yes	9	29.0
	No	22	71.0
Microtransit use	Commuting to work	21	67.7
	Leisure	17	54.8
	Daily errand	23	74.2
	Doctors' appointment	27	87.1
	Other	4	12.9
Microtransit frequency	2-3 times per week	11	36.7
	4-5 times per week	18	60.0
	Once per week	1	3.3
Microtransit travel times	Weekday early mornings (5:30am-9:30am)	17	54.8
	Weekday late mornings (9:30am-12 noon)	12	38.7
	Weekday afternoon (12 noon-4pm)	20	64.5
	Weekday afternoon/evening (4pm-7pm)	15	48.4
	Weekend mornings (7am-11:30am)	9	29.0
	Weekend afternoons (12 noon-6pm)	8	25.8

Table 8: ACT-R hyperparameters.

Parameter	Value
subsymbolic	True
utility noise	0.01
utility learning	True
utility alpha	0.01
strict harvesting	True

Table 9: Multiarmed bandit hyperparameters.

Parameter	Value
exploration rate	0.3
uniform distribution parameter	0.2
decay rate	0.8
reward weight	0.8
initial weight	0.5

Appendix B. Reproducibility Details

B.1 Hyperparameters

In our experiments, we used the hyperparameters specified in Tables 8 and 9 for our ACT-R agents and bandits respectively. For the CARS spatial tolerance learning agent, we experimented with the (default) PPO hyperparameters in Table 10 as well as with modified learning rate and number of steps (`n_steps`). A detailed sensitivity analysis result can be found in Appendix D, Table 25.

B.2 Rider Parameters

While running our experiments, we considered riders with the internal and shared parameters specified in Tables 11 through 22.

Table 10: Default PPO hyperparameters.

Parameter	Value
policy	MlpPolicy
learning rate	0.0003
number of steps	2048
batch size	64
number of epochs	10
verbose	1

Table 11: Competitive rider: Profile

Parameter	Value
Gender	Male
Age	Early 30s
Weather feature trace	Sunny > Rainy
Time of day feature trace	Morning = Afternoon = Evening

Table 12: Competitive rider: Spatial tolerance

Weather	Time of Day	Spatial Tolerance (in meters)
Sunny	Morning	300
Sunny	Afternoon	300
Sunny	Evening	300
Rainy	Morning	50
Rainy	Afternoon	50
Rainy	Evening	50

Table 13: Competitive rider: Empathy (persuasive value phrase) distribution

Value Phrase	Persuasive Percentage
Health benefit	100 %

Table 14: Individualistic rider: Profile

Parameter	Value
Gender	Female
Age	16
Weather feature trace	sunny > rainy
Time of day feature trace	[Morning = Afternoon] > Evening

Table 15: Individualistic rider: Spatial tolerance

Weather	Time of Day	Spatial Tolerance (in meters)
Sunny	Morning	450
Sunny	Afternoon	450
Sunny	Evening	20
Rainy	Morning	150
Rainy	Afternoon	150
Rainy	Evening	10

Table 16: Individualistic rider: Empathy (persuasive value phrase) distribution

Value Phrase	Persuasive Percentage
Ill people	40 %
Senior citizens	36 %
Environmental benefit	24 %

Table 17: Moderate rider: Profile

Parameter	Value
Gender	Female
Age	Mid 40s
Weather feature trace	Sunny > Rainy
Time of day feature trace	Morning > Afternoon > Evening

Table 18: Moderate rider: Spatial tolerance

Weather	Time of Day	Spatial Tolerance (in meters)
Sunny	Morning	475
Sunny	Afternoon	400
Sunny	Evening	50
Rainy	Morning	250
Rainy	Afternoon	200
Rainy	Evening	20

Table 19: Moderate rider: Empathy (persuasive value phrase) distribution

Value Phrase	Persuasive Percentage
Babies	40%
Health benefit	60%

Table 20: Altruistic rider: Profile

Parameter	Value
Gender	Female
Age	Early 70s
Weather feature trace	Sunny > Rainy
Time of day feature trace	Morning > Afternoon > Evening

Table 21: Altruistic rider: Spatial tolerance

Weather	Time of Day	Spatial Tolerance (in meters)
Sunny	Morning	350
Sunny	Afternoon	300
Sunny	Evening	50
Rainy	Morning	100
Rainy	Afternoon	100
Rainy	Evening	20

Table 22: Altruistic rider: Empathy (persuasive value phrase) distribution

Value Phrase	Persuasive Percentage
Babies	55%
Children	45%

Table 23: Average Accepted Adjustment and Acceptance Percentage for the three models and the baseline (untrained), measured over 500 episodes of 1,000 time steps each. SAT+SPATIAL outperforms the other models, getting riders to walk further and behave prosocially more frequently, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.

Experiment	Metric	SAT+SPATIAL	SPATIAL	PROFILE	UNTRAINED
Competitive	Accepted Adjustment	160.09	76.77	73.13	73.01
Competitive	Acceptance %	62.41	44.42	43.79	43.96
Individualistic	Accepted Adjustment	194.22	95.9	95.81	95.23
Individualistic	Acceptance %	63.51	47.14	47.21	47.15
Moderate	Accepted Adjustment	207.17	138.22	137.45	137.28
Moderate	Acceptance %	73.35	63.68	63.79	63.68
Prosocial	Accepted Adjustment	297.56	251.87	249.31	250.42
Prosocial	Acceptance %	99.94	99.94	99.94	99.94
Altruistic	Accepted Adjustment	299.45	251.03	248.55	250.16
Altruistic	Acceptance %	99.96	99.96	99.96	99.96

Table 24: Effect sizes measured as Cohen’s d scores and corresponding 95% confidence intervals of the improvement of SAT+SPATIAL over SPATIAL, PROFILE and an untrained model for (1) accepted adjustment value and (2) the percentage of accepted suggestions. These are measured over 500 episodes, each of length 1,000 time steps. SAT+SPATIAL outperforms the other models, getting riders to walk further and behave prosocially more frequently, especially for the less prosocial personas; this indicates that considering rider satisfaction is necessary to design a prosocial ecosystem.

Experiment	Metric	Vs SPATIAL	CI	Vs PROFILE	CI	Vs UNTRAINED	CI
Competitive	Accepted Adjustment	0.56	0.47 to 0.65	0.59	0.49 to 0.68	0.59	0.49 to 0.68
Competitive	Acceptance %	9.93	9.31 to 10.55	10.42	9.77 to 11.07	10.13	9.5 to 10.77
Individualistic	Accepted Adjustment	0.51	0.42 to 0.6	0.51	0.42 to 0.62	0.52	0.42 to 0.61
Individualistic	Acceptance %	8.9	8.34 to 9.46	9.01	8.44 to 9.57	8.99	8.42 to 9.55
Moderate	Accepted Adjustment	0.45	0.36 to 0.54	0.45	0.36 to 0.55	0.45	0.36 to 0.55
Moderate	Acceptance %	6.4	5.99 to 6.81	6.57	6.15 to 6.99	6.58	6.16 to 6.99
Prosocial	Accepted Adjustment	0.33	0.24 to 0.42	0.34	0.25 to 0.43	0.34	0.25 to 0.43
Prosocial	Acceptance %	0	-0.09 to 0.09	0	-0.09 to 0.09	0	-0.09 to 0.09
Altruistic	Accepted Adjustment	0.31	0.22 to 0.4	0.33	0.24 to 0.42	0.32	0.23 to 0.41
Altruistic	Acceptance %	0	-0.09 to 0.09	0	-0.09 to 0.09	0	-0.09 to 0.09

Appendix C. Additional Results

C.1 Spatial Tolerance Learning Results

Tables 23 and 24 show the performance of the three models and the untrained baseline over 500 episodes.

C.2 Persuasive Strategy Learning Results

Figures 12 through 15 show our persuasive strategy results.

Appendix D. Sensitivity Analysis




Figure 12: Competitive rider: actual vs top three predicted persuasive value phrase percentages; the remaining 1% is divided between the rest of the phrases.




Figure 13: Individualistic rider: actual vs top three predicted persuasive value phrase percentages; the remaining 3% is divided between the rest of the phrases.




Figure 14: Moderate rider: actual vs top three calculated persuasive value phrase percentages; the remaining 36% is divided between the rest of the phrases.

Table 25: Sensitivity analysis, keeping all other parameters constant and varying learning rate (lr) and number of steps (n_steps) for each model. The default model uses the parameters in Table 10. Results were measured over 500 episodes.

Experiment	Model	Metric	Default	lr=0.003	lr=0.00003	n_steps=1024	n_steps=4096
Competitive	SAT+SPATIAL	Accepted Adjustment	157.76	160.09	137.46	158.17	148.38
		Acceptance %	66.16	62.41	49.93	51.53	59.27
		Rider Satisfaction	8.12	7.76	6.06	6.38	7.05
	SPATIAL	Accepted Adjustment	76.77	176.6	73.29	76.34	76.33
		Acceptance %	44.42	50.03	43.95	44.56	44.47
		Rider Satisfaction	5.27	3.47	5.23	5.29	5.27
	PROFILE	Accepted Adjustment	73.13	75.69	72.95	72.97	73.02
		Acceptance %	43.79	42.28	43.89	43.82	43.83
		Rider Satisfaction	5.21	5.01	5.22	5.22	5.22
Individualistic	SAT+SPATIAL	Accepted Adjustment	191.58	194.22	123.57	202.62	184.24
		Acceptance %	63.48	63.51	40.13	65.93	62.02
		Rider Satisfaction	392.91	393.2	247.6	406.11	378.68
	SPATIAL	Accepted Adjustment	95.9	128.29	95.41	95.6	96.08
		Acceptance %	47.14	43.33	47.19	47.13	47.21
		Rider Satisfaction	290.18	257.79	290.59	290.18	290.62
	PROFILE	Accepted Adjustment	95.81	96.58	95.56	95.86	95.67
		Acceptance %	47.21	46.85	47.11	47.18	47.22
		Rider Satisfaction	290.68	288.25	290.14	290.46	290.72
Moderate	SAT+SPATIAL	Accepted Adjustment	176.24	207.17	137.84	177.68	173.09
		Acceptance %	61.13	73.35	63.5	62.02	62.56
		Rider Satisfaction	15096.7	18274.1	15 457.65	14 938.18	15 426.69
	SPATIAL	Accepted Adjustment	138.22	183.7	137.71	138.15	138.5
		Acceptance %	63.68	63.28	63.76	63.62	63.68
		Rider Satisfaction	15488.85	12979.1	15 509.49	15 462.89	15 475.38
	PROFILE	Accepted Adjustment	137.45	137.07	137.21	137.81	137.38
		Acceptance %	63.79	64.92	63.64	63.75	63.67
		Rider Satisfaction	15534.06	16132.88	15 499.37	15 521.32	15 524.65
Prosocial	SAT+SPATIAL	Accepted Adjustment	277.28	297.56	251.42	270.05	274.83
		Acceptance %	99.94	99.94	99.94	99.94	99.94
		Rider Satisfaction	38099.45	44448.36	38 885.76	38 086.88	37 579.15
	SPATIAL	Accepted Adjustment	251.87	413.1	250.78	253.34	254.09
		Acceptance %	99.94	99.94	99.94	99.94	99.94
		Rider Satisfaction	38820.81	23390.01	38 972.97	38 771.27	38 696.8
	PROFILE	Accepted Adjustment	249.31	228.32	250.3	249.42	248.93
		Acceptance %	99.94	99.94	99.94	99.94	99.94
		Rider Satisfaction	39119.61	41952.72	39 067.81	39 182.41	39 164.04
Altruistic	SAT+SPATIAL	Accepted Adjustment	249.41	299.45	250.32	249.23	251.5
		Acceptance %	99.96	99.96	99.96	99.96	99.96
		Rider Satisfaction	49549.15	50662.33	50 361.31	49 927.83	49 417.35
	SPATIAL	Accepted Adjustment	251.03	372.42	250.7	253.58	253.9
		Acceptance %	99.96	99.96	99.96	99.96	99.96
		Rider Satisfaction	49980.35	30123.69	50 134.23	49 625.71	49 473.6
	PROFILE	Accepted Adjustment	248.55	187.97	250.69	248.34	248.35
		Acceptance %	99.96	99.96	99.96	99.96	99.96
		Rider Satisfaction	50447.23	57257.95	50 089.8	50 472.88	50 448.69




Figure 15: Altruistic rider: actual vs top three predicted persuasive value phrase percentages; the remaining 4% is divided between the rest of the phrases.