Integrating neural network-based habit learning into random utility models for transport mode choice

Yue Yang[1] and Elisabetta Cherchi[1]

[1]Division of Engineering, New York University Abu Dhabi, United Arab Emirates

Short summary

Discrete choice models are very powerful models able to incorporate a large range of behaviours. However, the theoretical interpretability that is a key feature of these models often limits the possibility of adequately modelling the behavioural complexity behind the decision process, despite the availability of a rich dataset. An exemplary case of such complexity is the inertia effect. Several functional forms are used in the literature, each based on a specific type of information. This study exploits the potential of Neural Networks to autonomously extract the temporal dynamics of habit, leading to inertia. Specifically, the network extracts a continuous 'history-support' index that adapts dynamically to the decision environment. These effects are subsequently integrated and tested within a standard, interpretable random utility model, and the results are compared with the same effects measured in the traditional way. In an empirical application using the six-week Mobidrive panel from Karlsruhe, Germany, the results confirm that the habit component exhibits a large individual effect and significantly improves the explanation of observed choices. More interestingly, it reveals that the learned habit influence exhibits temporal decay and varies significantly by trip purpose, capturing rich behavioural dynamics that traditional static frequency measures fundamentally fail to represent.

Keywords: neural network, discrete choice analysis; habit panel data; mode choice

1 Introduction

For decades, discrete choice models have served as the primary analytical framework for understanding and predicting these decisions. These models provide structural interpretability and produce theoretically consistent and policy-relevant parameters that capture key behavioural trade-offs, as well as richer behavioural effects such as the impact of habit, among others. The theoretical interpretability embedded in DCM represents a key feature of these models, but at the same time, the operationalization of behavioural theory often limits the possibility of adequately modelling the behavioural complexity underlying the decision process, despite the availability of rich datasets. As travel behaviour data have become richer in detail and structure, there has been growing interest in flexible, data-driven methods. Within the mode choice context, many studies have applied machine learning approaches, such as deep neural networks, as standalone predictive tools for travel mode choice (Omrani, 2015; F. Wang & Ross, 2018; Chen & Cheng, 2023; Yang et al., 2024; Mo et al., 2023; Zhang & Xu, 2026). While such data-driven methods can improve predictive performance, their limited transparency makes it difficult to recover the behavioural structure required for economic interpretation, welfare analysis, and forecasting under counterfactual policy scenarios (Van Cranenburgh et al., 2022).

A more promising direction is to integrate data-driven methods and discrete choice models within a single behavioural framework, where flexible machine learning models can be used to capture behavioural patterns that are difficult to represent within conventional utility specifications, while the random utility framework is retained to preserve interpretability and policy relevance. In this line, much of the existing literature focuses primarily on learning non-linear effects from observed cross-sectional variables, such as socio-demographic characteristics and trip attributes. For example, models such as the Learning Multinomial Logit (Sifringer et al., 2020) and TasteNet (Han et al., 2022) use neural networks to capture unobserved taste heterogeneity by learning how sensitivities to time or cost vary across individuals as a function of socio-demographic characteristics. Similarly, models such as ResLogit (Wong & Farooq, 2021) and TB-ResNet (S. Wang et al., 2021) employ

neural architectures to uncover complex, non-linear trade-offs based on static trip attributes, with the aim of improving how the utility of an alternative is derived from its physical attributes.
In this paper, we would like to leverage data-driven methods to explicitly measure habit formation and integrate it into a standard random utility model. Conventional discrete choice models have accounted for habit leading to inertia by using lagged variables, or the simple frequency of past trip repetitions, or by specifying formulations that incorporate the previously experienced utility of prior choices over predefined memory windows. Most specifications use a short window: i.e., the previous trip or the trip performed before an event (Cantillo et al., 2007; González et al., 2017; La Paix et al., 2022; Yáñez et al., 2025). Few papers have used panel data that allow for longer memory windows (Cherchi & Cirillo, 2014), but here the sequence of past trips is collapsed into a single summary measure, a frequency count. As a result, the temporal structure of the data is largely discarded, and a history of choices is reduced to a set of unordered observations rather than treated as an evolving process of reinforcement and decay. To try to capture these dynamics, we use sequential deep learning methods, such as recurrent neural networks. Past decisions are processed as an ordered sequence and maintain an internal memory state, allowing us to more naturally represent the accumulation and fading of habit over time. In this paper, we combine this learning-based representation of travellers' habits with a standard attribute-based random utility specification, preserving the interpretability and policy relevance of the discrete choice framework.

2 Methodology

Habit representation from panel history

The basic idea for capturing inertia more flexibly consists of treating the finite look-back window of past decisions as an ordered sequence of experiences. For each current choice occasion, the model examines the last T observed trips, preserves their chronological order, and constructs a context-dependent summary of those past experiences. As illustrated in Figure 1, the procedure has three steps: each past trip is first represented numerically; the sequence of the last T trips is then processed in chronological order; and finally, the model assigns larger or smaller weights to different past trips according to how relevant they are for the current trip. Behaviourally, the resulting habit term is therefore not a simple frequency count or average over the last T trips, but a context-dependent summary of recent experience.
Operationally, for each current decision occasion t, the network takes as input the traveller's most recent T observed trips and the current trip context, and returns an alternative-specific habit vector $\mathbf{H}_{n,t}$. As an empirical example, if $T = 10$, the habit score for trip 11 is constructed from trips 1–10, the habit score for trip 12 is constructed from trips 2–11, and so forth. The same network, with one common set of parameters, is used for all individuals and all choice occasions. These neural-network parameters are estimated jointly with the other parameters (i.e., from attribute-based utility and the LLM-derived choice-set prior) by maximizing the simulated log-likelihood over all observations from trip 11 onward. As in many dynamic discrete choice formulations, the history window rolls forward one trip at a time. The difference here is not the use of a rolling window itself, but that the order of trips within the window is retained and the weighting of past trips is learned from the data rather than fixed *a priori*. Thus, the history window changes from one trip to the next, but the network itself is not separately re-estimated for each traveller or each new trip. Importantly, this does not imply that the resulting weights are static. The estimated network parameters are common across the sample, but the realised habit scores and attention weights vary across individuals and decision occasions because they are functions of the observed history and the current trip context. Because the estimation sample starts from trip 11, each modelled choice is associated with a complete observed history of the previous $T = 10$ trips, so no artificial zero-padding or masking is required in the reported estimation.
Let $\mathbf{h}_{n,\tau}$ denote the observed state of individual n at a past decision occasion τ, including the chosen mode and relevant trip descriptors (such as experienced travel time and cost). In the empirical application, this past-history input may include the previously chosen mode, the experienced travel time and travel cost of that chosen mode, the past trip purpose, and the past day of the week. Each past state is first mapped into a numerical representation that captures the information of the past experience,

$$\mathbf{z}_{n,\tau} = \mathrm{Emb}(\mathbf{h}_{n,\tau}), \tag{1}$$

where $\mathrm{Emb}(\cdot)$ denotes a trainable linear embedding function. Equation 1 is applied to each of the T past trips in the look-back window, that is, for $\tau = t - T, \ldots, t - 1$, rather than only

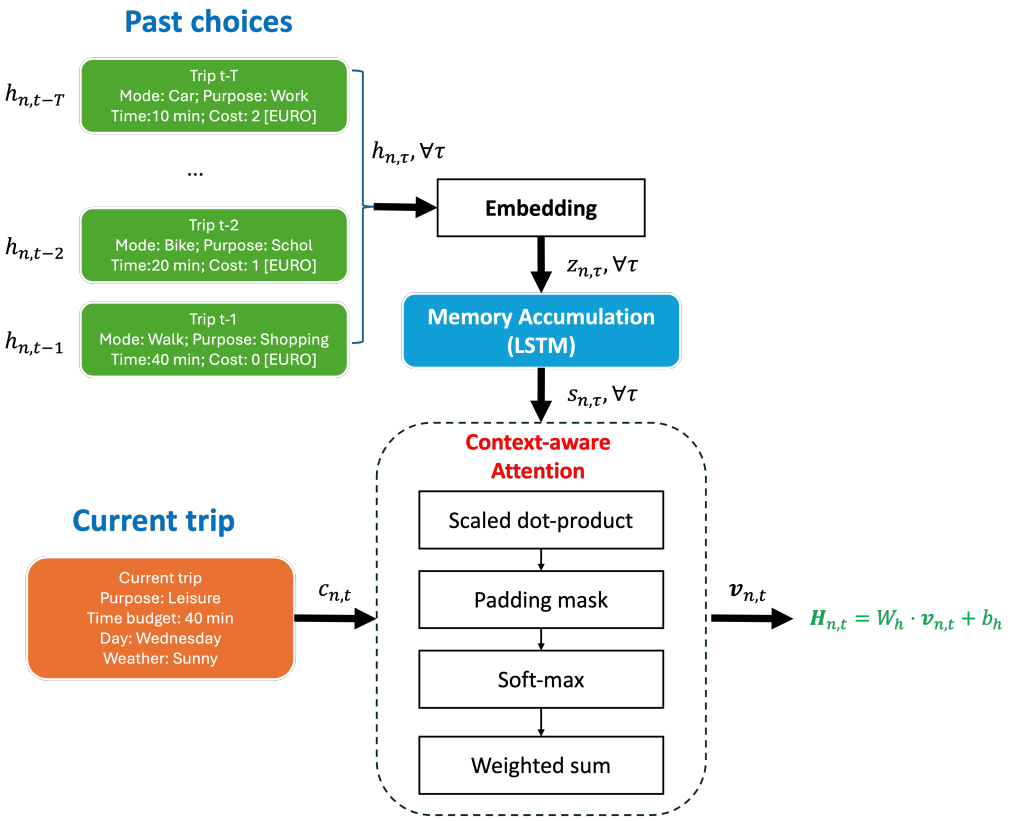


Figure 1: Architecture of the dynamic habit learning framework.

to the immediately preceding trip. The embedded history states are then processed sequentially through a Long Short-Term Memory (LSTM) network, which builds a running summary of past travel experience. In particular, the LSTM reads the T past trips in chronological order, so that the representation at each step retains information from the trips that came before it. In this way, the sequence of experiences is preserved rather than collapsed into a single unordered summary measure. Rather than requiring the analyst to specify in advance how the influence of the first, second, ..., T-th lag should vary within the chosen window, the model learns from the data how information from earlier and more recent trips should be carried forward in the history representation. Through this memory accumulation process, past choices within the look-back window of length T are summarized into a sequence of hidden states:

$$\mathbf{s}_{n,\tau} = \mathrm{LSTM}(\mathbf{z}_{n,\tau}, \mathbf{s}_{n,\tau-1}). \tag{2}$$

Because the relevance of a past trip depends on the current decision being predicted, the model next applies a context-dependent weighting step, referred to here as an attention mechanism. Let $\mathbf{c}_{n,t}$ denote the current trip context vector, including variables such as trip purpose and time budget. In the empirical application, this current-context input could include the current trip purpose, the day of the week, the panel week number, the available time budget, whether the traveller has children, and whether the current tour involves intermediate stops. For each past occasion τ, a raw relevance score is computed as

$$e_{n,\tau,t} = \frac{(W_q \mathbf{c}_{n,t})^\top (W_k \mathbf{s}_{n,\tau})}{\sqrt{d}}, \tag{3}$$

where W_q and W_k are trainable weight matrices that project the context query and historical keys into a shared space of dimension d, and $\sqrt{d}$ acts as the scaling factor. Equation 3 compares the current trip context with each historical state $\mathbf{s}_{n,\tau}$, rather than with each raw past trip in isolation. Since each $\mathbf{s}_{n,\tau}$ already summarizes the ordered history up to past occasion τ, the resulting score indicates how behaviourally relevant that past experience is for the present decision.
The relevance scores are then converted into attention weights through the softmax transformation:

$$\omega_{n,\tau,t} = \frac{\exp(e_{n,\tau,t})}{\sum_{u=t-T}^{t-1} \exp(e_{n,u,t})}. \tag{4}$$

These weights are non-negative and sum to one, so they can be interpreted as the share of influence assigned to each of the last T trips. Accordingly, Eq. 4 assigns a larger weight to past experiences that are more relevant to the current context and a smaller weight to those that are less relevant. This relevance is not defined by a single manually specified criterion, such as sharing the same mode or the same purpose. Instead, it is learned jointly from the variables entering the past-history input $\mathbf{h}_{n,\tau}$ and the current-context input $\mathbf{c}_{n,t}$, so it can depend on combinations of past chosen mode, experienced time and cost, past purpose, past day of week, and current trip characteristics. Consequently, the same recent history can yield different weight patterns when the current trip context changes.
The resulting context-aware history summary is

$$\mathbf{v}_{n,t} = \sum_{\tau=t-T}^{t-1} \omega_{n,\tau,t} \mathbf{s}_{n,\tau}, \tag{5}$$

which is a weighted sum of the past hidden states, with more relevant trips receiving larger weights. This summary is then mapped into an alternative-specific habit vector through a linear projection,

$$\mathbf{H}_{n,t} = W_h \mathbf{v}_{n,t} + b_h, \tag{6}$$

where the i-th element, $H_{in,t}$, represents the habitual score of alternative i for individual n at occasion t. Although this output is alternative-specific, it is not independent of the current purpose or timing of the trip. Because the attention weights depend on the current context vector $\mathbf{c}_{n,t}$, the same traveller can obtain different habit vectors from the same recent history when predicting, for example, a work trip versus a leisure trip, or one day of the week versus another.
This habit score represents a history-support index for each mode and captures how strongly the recent sequence of past choices supports choosing that mode again in the current trip context. In this sense, it is not only or exactly a measure of repetition, but it also reflects whether past choices were recent, stable, or instead involved frequent switching across modes. A positive score means that the recent history is more supportive of choosing that mode now. A negative score does not mean "negative habit" in itself, but rather that the recent history provides weaker support for that mode, or is more consistent with other alternatives.
In the empirical application with five travel modes (Car Driver, Car Passenger, Public Transit, Walking, and Cycling), this output takes the form of a 5×1 column vector:

$$\mathbf{H}_{n,t} = \begin{bmatrix} H_{\mathrm{CD}} \\ H_{\mathrm{CP}} \\ H_{\mathrm{PT}} \\ H_{\mathrm{W}} \\ H_{\mathrm{B}} \end{bmatrix}.$$

These continuous values directly quantify the context-adjusted behavioural inertia toward each respective mode. Eventually, this habit vector enters the joint utility linearly as specified in Eq. 7.

Integrated Utility Specification

Having defined the neural-network-based habit representation, the next step is to incorporate it into the discrete choice model. Figure 2 provides an overview of the proposed modelling framework. The framework builds on a standard random utility model for travel mode choice, into which the habit term is added to reflect the influence of past behaviour on the probability of choosing each mode. In this way, the neural-network-based habit component does not replace the discrete choice structure, but complements it by providing a flexible, data-driven representation of behavioural persistence. Unlike conventional discrete choice specifications, where inertia is represented through predefined parametric terms, the present framework uses the learned habit representation from observed travel history and integrates it directly into the utility. The neural-network component and the discrete choice model are estimated jointly within a single optimization problem.
For individual n, alternative $i \in \mathcal{C}_n$, and decision occasion t, the proposed framework specifies utility as the sum of two components and a stochastic term:

$$U_{in,t} = \underbrace{V_{in,t}^{\mathrm{arr}}}_{\text{attribute-based utility}} + \underbrace{\alpha_i H_{in,t}}_{\text{habit term}} + \varepsilon_{in,t}, \tag{7}$$

where $V_{in,t}^{\mathrm{arr}}$ denotes the attribute-based utility derived from observed trip and individual characteristics (which explicitly incorporates crossed random effects for travel time and cost sensitivities),

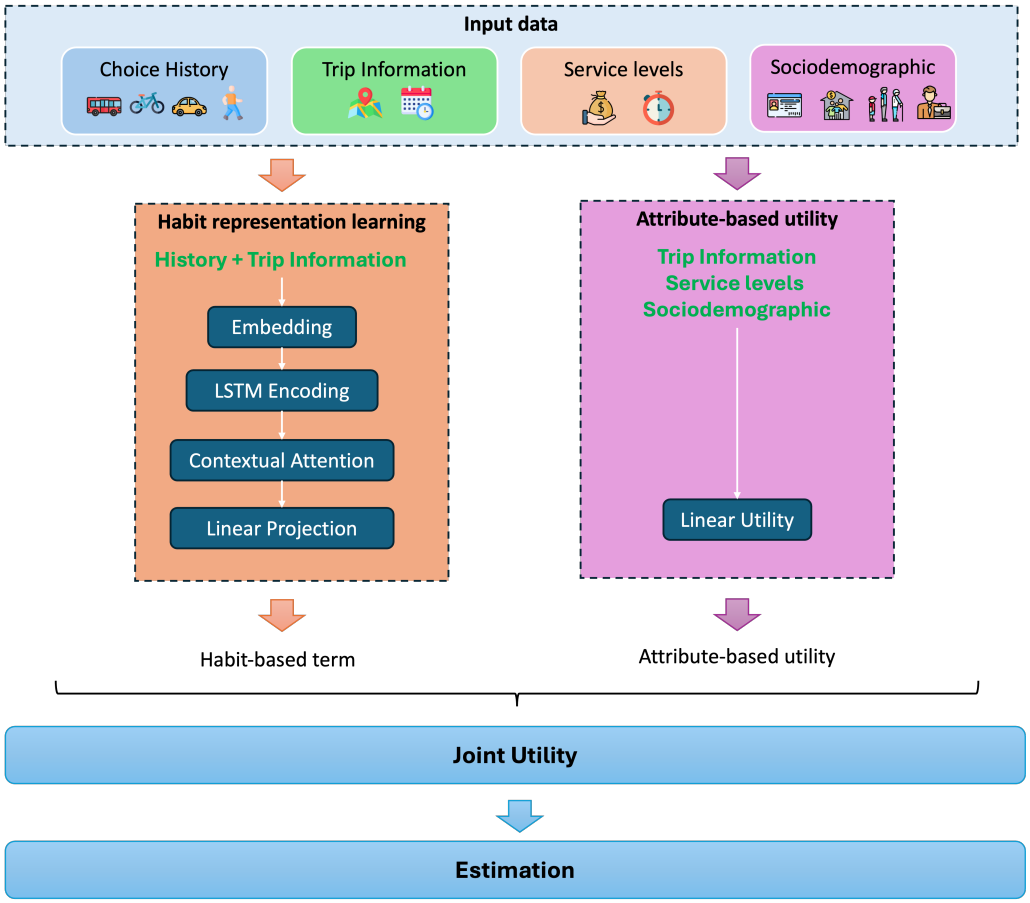


Figure 2: Overview of the proposed dynamic travel mode choice framework.

and $H_{in,t}$ denotes the habit score derived from historical choices scaled by an alternative-specific behavioural parameter α_i.

The attribute-based utility captures the conventional economic trade-offs associated with observed service levels and socio-demographic characteristics. For individual n, alternative i, and decision occasion t, the attribute-based utility is specified as:

$$\begin{aligned} V_{in,t}^{\text{arr}} = & \, ASC_i \\ & + \beta_{\text{time},n,t}\, x_{\text{time},in,t} + \beta_{\text{cost},n,t}\, x_{\text{cost},in,t} \\ & + \sum_{k=1}^{K} \beta_{ik} x_{in,t,k} \\ & + \sum_{m=1}^{M} \delta_{im} S_{n,m} \\ & + \sum_{k=1}^{K} \sum_{p=1}^{P} \theta_{ikp} x_{in,t,k} Z_{n,t,p}, \end{aligned} \tag{8}$$

where ASC_i is the alternative-specific constant for mode i, $x_{\text{time},in,t}$ and $x_{\text{cost},in,t}$ denote travel time and cost respectively, $x_{in,t,k}$ denotes the remaining k-th alternative-specific attributes, $S_{n,m}$ denotes the m-th socio-demographic characteristic of individual n, and $Z_{n,t,p}$ denotes the p-th trip-specific contextual variable. The parameters β_{ik}, δ_{im}, and θ_{ikp} measure, respectively, the direct effects of service attributes, the influence of socio-demographic characteristics, and the interaction effects between service attributes and decision context. The interaction term $x_{in,t,k} Z_{n,t,p}$ allows the effect of service attributes to vary across trip situations. For identification, one alternative-specific constant is normalized to zero.

Following the specification in Cherchi et al. (2017), the random coefficients for travel time and cost are formulated using a crossed error structure to capture variability at both the weekly and daily

levels:

$$\beta_{\text{time},n,t} = \bar{\beta}_{\text{time}} + \sigma_{\text{time},w}\,\xi^{\text{time}}_{n,w(t)} + \sigma_{\text{time},d}\,\xi^{\text{time}}_{n,d(t)}, \tag{9}$$

$$\beta_{\text{cost},n,t} = \bar{\beta}_{\text{cost}} + \sigma_{\text{cost},w}\,\xi^{\text{cost}}_{n,w(t)} + \sigma_{\text{cost},d}\,\xi^{\text{cost}}_{n,d(t)}, \tag{10}$$

where $\bar{\beta}_{\text{time}}$ and $\bar{\beta}_{\text{cost}}$ are the population mean sensitivities. The mappings $w(t)$ and $d(t)$ return the specific week and day of the week corresponding to decision occasion t. The terms $\xi \sim N(0,1)$ are independent standard normal draws, while $\sigma_{\text{time},w}$, $\sigma_{\text{time},d}$, $\sigma_{\text{cost},w}$, and $\sigma_{\text{cost},d}$ are the standard deviations to be estimated. This structure explicitly accounts for correlation across choices made by the same individual within the same week or on the same day of the week.
This error structure yields a crossed-panel mixed logit model that explicitly accounts for day-to-day and week-to-week correlation across repeated observations from the same individual. The term $\varepsilon_{in,t}$ is an i.i.d. Gumbel error that captures remaining unobserved variation across decision occasions.

Estimation

The proposed framework is estimated jointly so that the three components of the utility specification are calibrated within a single optimization problem. Let

$$\Theta = \left\{ \bar{\beta}_{\text{time}}, \bar{\beta}_{\text{cost}}, \sigma_{\text{time},w}, \sigma_{\text{time},d}, \sigma_{\text{cost},w}, \sigma_{\text{cost},d}, \beta_{ik}, \delta_{im}, \theta_{ikp}, ASC_i, \alpha, \Phi_H \right\}$$

denote the full parameter set, where the σ terms denote the standard deviations of the crossed random effects, Φ_H collects the trainable parameters of the neural-network-based habit component, and w_l denotes the aggregation weights used to combine the outputs of the L large language model agents.
Let $\boldsymbol{\xi}_n$ denote the vector collecting all independent standard normal draws $(\xi^{\text{time}}_w, \xi^{\text{time}}_d, \xi^{\text{cost}}_w, \xi^{\text{cost}}_d)$ for individual n across all their respective weeks and days observed in the panel. The choice probability conditional on a realization of $\boldsymbol{\xi}_n$ takes the logit form:

$$L(y_{n,t} = i \mid \boldsymbol{\xi}_n, \mathcal{F}_{n,t}; \Theta) = \frac{\exp\big(V^{\text{arr}}_{in,t}(\boldsymbol{\xi}_n) + \alpha H_{in,t} + \gamma \ln(\tilde{q}_{in,t})\big)}{\sum_{j \in \mathcal{C}_n} \exp\big(V^{\text{arr}}_{jn,t}(\boldsymbol{\xi}_n) + \alpha H_{jn,t} + \gamma \ln(\tilde{q}_{jn,t})\big)}. \tag{11}$$

Since the random coefficients map deterministically to specific decision occasions based on the day and week, the unconditional probability of the observed choice sequence for individual n is obtained by integrating over the standard multivariate normal distribution of the draws:

$$P(\mathbf{y}_n \mid \mathcal{F}_n; \Theta) = \int \prod_t L(y_{n,t} \mid \boldsymbol{\xi}_n, \mathcal{F}_{n,t}; \Theta)\, \phi(\boldsymbol{\xi}_n \mid \mathbf{0}, \mathbf{I})\, d\boldsymbol{\xi}_n, \tag{12}$$

where $\phi(\boldsymbol{\xi}_n \mid \mathbf{0}, \mathbf{I})$ denotes the multivariate standard normal density. This integral does not have a closed-form solution and is approximated by simulation:

$$\hat{P}(\mathbf{y}_n \mid \mathcal{F}_n; \Theta) = \frac{1}{R} \sum_{r=1}^{R} \prod_t L(y_{n,t} \mid \boldsymbol{\xi}^{(r)}_n, \mathcal{F}_{n,t}; \Theta), \tag{13}$$

where $\boldsymbol{\xi}^{(r)}_n$ denotes the r-th set of draws and R is the number of simulation draws. The model is estimated by maximizing the simulated log-likelihood:

$$\mathcal{SL}(\Theta) = \sum_n \log \hat{P}(\mathbf{y}_n \mid \mathcal{F}_n; \Theta). \tag{14}$$

The model is estimated end-to-end by maximizing the simulated log-likelihood with respect to the full parameter set Θ. Accordingly, the neural-network parameters in Φ_H, the LLM aggregation weights w_l, the standard deviations of the crossed panel effects, and the coefficients entering the utility specification in Eq. 8 are all estimated jointly within the same optimization procedure. The product over decision occasions inside the integral in Eq. 12 ensures that the panel structure is respected: the draws for a specific week or day are applied consistently to all trips made by the individual within that time frame, inducing the desired temporal correlation.

3 Results

The proposed framework is empirically evaluated using data from the Mobidrive continuous travel survey, conducted in Karlsruhe, Germany, in 1999 (PTV AG et al., 2000). This survey provides a unique six-week longitudinal panel that tracks the continuous daily travel diaries of participants, making it an ideal empirical setting to study behavioural persistence and contextual mode switching. To capture consistent day-to-day and week-to-week variability in routine travel, the sample is restricted to the five working days of each week, explicitly excluding weekends (Saturdays and Sundays) since weekend activity patterns and constraints differ substantially from the typical work week. The final dataset used in this paper comes from Cherchi & Cirillo (2014); Cherchi et al. (2017), and contains 4,089 tours made by 129 individuals from 56 households. For each tour, the available mode choice set consists of five alternatives: car driver, car passenger, public transport, walking, and biking. Figure 3 reports the aggregate distribution of observed choices across the sample, showing that overall mode shares remain remarkably stable across the observation period. However, this aggregate stability masks complex intra-personal variability. To visually illustrate individual choice trajectories over time, Figure 4 displays the day-to-day mode choice sequences for a selected subset of ten individuals. This micro-level view clearly reveals a mix of rigid behavioural inertia (consistent use of a single mode) and occasional context-driven switching, precisely the dynamic mechanisms our integrated habit learning framework is designed to capture.

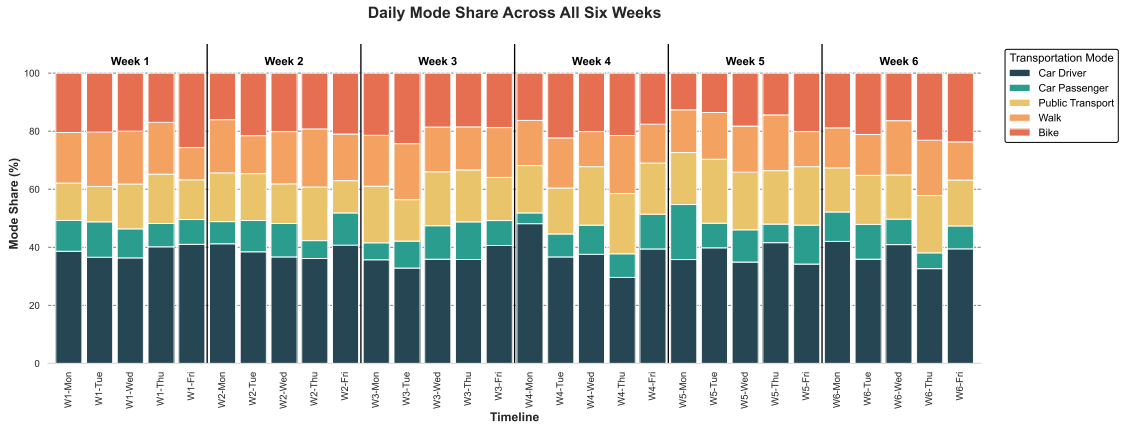


Figure 3: Aggregate daily mode share distribution for weekday tours across the six-week observation period.

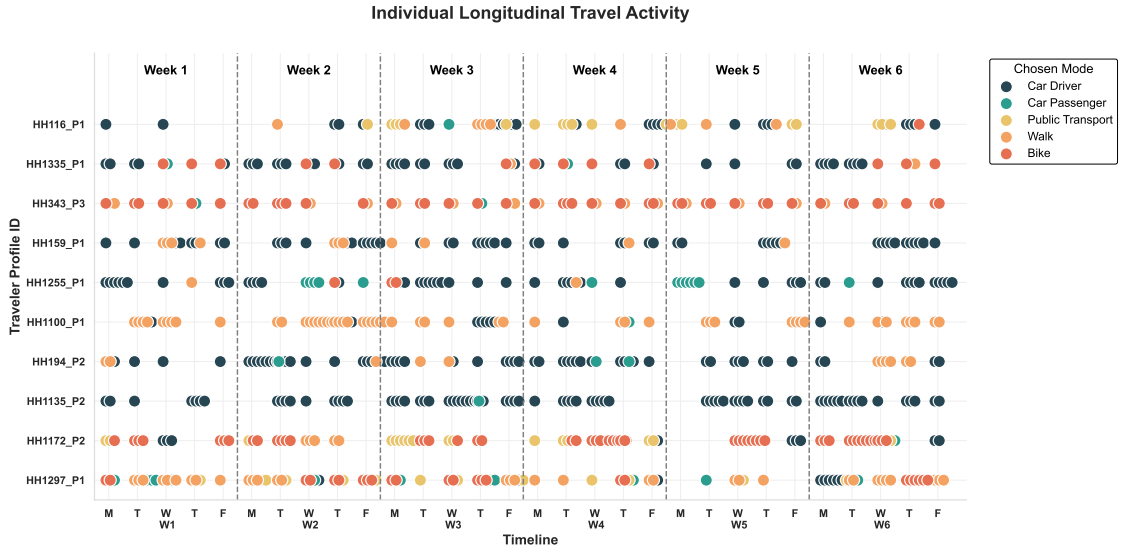


Figure 4: Individual longitudinal mode choice sequences over 30 continuous weekdays.

Model Specification

Table 1 summarizes the variables included in the attribute-based utility component, which serves as the conventional core of our choice model. This specification strictly adopts the mixed logit formulation proposed by Cherchi et al. (2017). The attribute-based utility captures standard behavioural trade-offs through four variable categories: alternative-specific constants, level-of-service attributes (travel time and cost), trip information, and socio-demographic characteristics. In particular, unobserved temporal heterogeneity is accommodated by estimating the sensitivities to travel time and cost as crossed random parameters with day- and week-level standard deviations. Meanwhile, observed heterogeneity is captured through systematic interaction terms between travel time, situational trip purposes, and individual characteristics.

Table 2 details the input variables for the dynamic habit learning component. To capture behavioural persistence, the neural network processes two inputs: a historical sequence of past experiences ($\mathbf{h}_{n,\tau}$) and the current trip context ($\mathbf{c}_{n,t}$), which acts as a situational filter to retrieve relevant habits. The memory window spans $T = 10$ previous trips. Each historical step includes the chosen mode, experienced time, cost, and past context, while the current context captures immediate constraints like trip purpose, day of the week, and time budget. The network's architecture (including the LSTM and attention mechanism) uses a latent dimension of $d = 16$ to synthesize complex, non-linear interactions between past behaviour and current needs. To address the "initial conditions" problem in panel data econometrics, each traveller's first 10 observed trips act exclusively as a historical buffer, with estimation starting from the 11th trip. This ensures a complete, empirically observed $T = 10$ sequence for every choice scenario, eliminating artificial data padding. Finally, the network consolidates this history into a 5×1 continuous habit vector ($\mathbf{H}_{n,t}$), which enters the baseline utility function directly as dynamic, alternative-specific habit inertia scores.

Table 1: Data dictionary of variables entering the attribute-based utility specification.

Category	**Variable Name**	**Description**
Model constants	ASC_i	Alternative-specific constant for mode i.
Level of service	Travel time	Mode-specific travel time (minutes).
	Travel cost	Out-of-pocket travel cost (DM).
	Time × Children	Travel time for married with children < 18.
	Time × Work	Travel time for work tours.
	Time × Education	Travel time for education tours.
	Time × Suburban (PT)	Travel time for suburban residents.
	Time × Age 51–65 (Bike)	Travel time for individuals aged 51–65.
Trip information	Leisure purpose (CP, PT)	1 if tour purpose is leisure.
	Work purpose (PT)	1 if tour purpose is work.
Socio-demographics	Age 18–25 (Bike)	1 if individual is aged 18–25.
	Age 26–35 (CD, Bike)	1 if individual is aged 26–35.
	Age 51–65 (PT, Walk, Bike)	1 if individual is aged 51–65.
	Main car user (CD)	1 if individual is the primary car user.
	Annual VMT (CD)	Annual vehicle mileage.
	Time budget (CD, Bike)	Available daily time budget (minutes).
	Urban residence (PT)	1 if household is in an urban area.
	Children (PT)	1 if married with children < 18.
	Season ticket (PT)	1 if individual holds a PT season ticket.

Evaluation

To evaluate the proposed methodology, we compare our integrated framework against two baseline models. This ablation study is designed to isolate the individual and joint contributions of the dynamic neural network (NN) habit learning component. To ensure a strictly fair comparison across all specifications, a uniform data treatment is applied: the first 10 observed trips for each individual serve exclusively as a historical buffer to initialize the habit and frequency variables. Consequently, log-likelihood maximization for all models commences strictly from the 11th trip onward.

Table 2: Input features for the dynamic habit learning neural network.

Input Type	Feature Name	Description
Historical sequence ($\mathbf{h}_{n,\tau}$)	Past chosen mode	One-hot encoded vector of the previously chosen alternative.
	Experienced time	Travel time of the previously chosen alternative.
	Experienced cost	Travel cost of the previously chosen alternative.
	Past trip purpose	Indicators for work and leisure purposes.
	Past day of week	Normalized day of the week of the past trip.
Current context query ($\mathbf{c}_{n,t}$)	Current purpose	Indicators for work, leisure, and education purposes.
	Temporal indicators	Normalized day of the week and panel week number.
	Time budget	Available daily time budget (scaled).
	Children	1 if married with children < 18.
	Stops	1 if the current tour involves intermediate stops.

We define the following three specifications for our experimental comparison, each characterized by a distinct formulation of the utility function $U_{in,t}$:

- **Model 1: Baseline Mixed Logit.** This model represents the traditional econometric approach, fully consistent with the specification proposed by Cherchi et al. (2017). It assumes no behavioural persistence (meaning no habit variables are included). The utility includes only the attribute-based component $V_{in,t}^{\text{attr}}$ in Equation 8.
- **Model 2: Mixed Logit (Frequency Habit).** This specification builds upon Model 1 by incorporating behavioural persistence via cumulative frequency variables. For any decision occasion $t \geq 11$, the frequency term $\text{Freq}_{in,t}$ counts the total number of previous trips made with a given alternative i up to trip $t-1$. To capture mode-specific inertia accurately, a distinct frequency coefficient $\beta_{\text{freq},i}$ is estimated for each travel mode. The utility function expands to incorporate this static habit measure:

$$U_{in,t}^{(2)} = V_{in,t}^{\text{attr}} + \beta_{\text{freq},i} \cdot \text{Freq}_{in,t} + \varepsilon_{in,t} \tag{15}$$

- **Model 3: Mixed Logit (NN Habit).** This model constitutes our fully integrated formulation in Equation 7. The neural network computes dynamic, alternative-specific inertia scores $H_{in,t}$ by updating the hidden state with all prior trips up to the current occasion. Comparing Model 3 against Model 2 explicitly isolates the performance gains achieved by modelling complex habit formation via representation learning rather than static frequency counts.

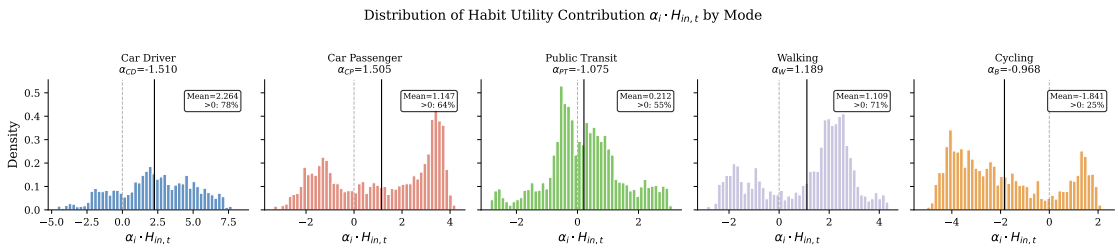


Figure 5: Distribution of the realised habit utility contribution ($\alpha_i H_{in,t}$) by travel mode.

Table 3 presents the estimation results. Model fit, evaluated via simulated log-likelihood, improves substantially from the baseline Mixed Logit (Model 1, $LL = -2061.68$) to the frequency-based specification (Model 2, $LL = -1589.44$), confirming strong behavioural inertia. However, substituting static counts with the dynamic neural network habit component (Model 3) yields a profound further improvement ($LL = -1133.03$). This confirms that treating past choices as an ordered, context-dependent sequence captures habit formation far more effectively than unordered frequency measures.

Crucially, core economic trade-offs remain structurally robust in the integrated model. Travel time and cost sensitivities remain highly significant, and interaction terms (e.g., Time × Education) remain stable. In contrast, variables proxying habitual patterns (`AnnualVMT_CD` and `SeasonTkt_PT`)

Table 3: Estimation Results for the Panel Mixed Logit Models

Variable	M1 (Baseline)	M2 (+Freq Habit)	M3 (+NN Habit)
Alternative Specific Constants			
ASC_CP	0.1378 (0.43)	0.9369 (2.27)	0.4329 (1.06)
ASC_PT	0.1719 (0.37)	0.1692 (0.28)	0.4273 (0.69)
ASC_W	0.7443 (2.08)	1.0534 (2.37)	0.1666 (0.36)
ASC_B	-1.1846 (-3.40)	-1.3064 (-2.65)	-0.6362 (-1.22)
Random Effects (Mean)			
β_{Time}	-0.0460 (-7.23)	-0.0649 (-8.27)	-0.0718 (-9.29)
β_{Cost}	-0.2151 (-4.26)	-0.2204 (-5.29)	-0.2201 (-5.58)
Random Effects (Standard Deviation)			
$\sigma_{\text{Time_week}}$	0.0611 (-34.02)	0.0645 (-32.91)	0.0619 (-29.65)
$\sigma_{\text{Time_day}}$	0.0654 (-22.66)	0.0677 (-30.99)	0.0695 (-28.91)
$\sigma_{\text{Cost_week}}$	0.1640 (-4.25)	0.1492 (n/a)	0.2120 (-10.71)
$\sigma_{\text{Cost_day}}$	0.1693 (-11.96)	0.1818 (-27.68)	0.1785 (-5.14)
Fixed Attributes			
Time × Children	-0.0613 (-5.88)	-0.0487 (-4.42)	-0.0676 (-4.93)
Time × Work	-0.0296 (-2.95)	-0.0046 (-0.47)	-0.0017 (-0.15)
Time × Education	0.0598 (5.69)	0.0792 (6.18)	0.0958 (6.08)
Time × Suburban_PT	0.0347 (4.40)	0.0302 (2.90)	0.0324 (2.62)
Time × Age65_B	-0.0384 (-4.14)	-0.0470 (-3.02)	-0.0422 (-2.62)
Leisure_CP	1.2856 (5.90)	1.4026 (5.67)	1.4367 (5.65)
Leisure_PT	-0.6863 (-2.74)	-0.2906 (-1.01)	-0.0380 (-0.12)
Work_PT	1.7384 (8.02)	1.1844 (4.74)	1.1546 (3.59)
Age18-25_B	1.0993 (3.86)	-0.2440 (-0.47)	1.2996 (2.86)
Age26-35_CD	2.7889 (5.04)	2.4041 (2.75)	1.7099 (1.93)
Age26-35_B	2.0501 (3.62)	1.5228 (2.28)	1.5114 (1.56)
Age51-65_PT	-0.2690 (-1.21)	0.2062 (0.83)	-0.2475 (-0.83)
Age51-65_W	1.2235 (3.27)	0.8846 (2.65)	1.2346 (2.97)
Age51-65_B	1.3964 (4.42)	1.0791 (2.51)	1.8213 (3.77)
MainCar_CD	1.7968 (9.50)	0.9989 (4.11)	1.2366 (5.28)
AnnualVMT_CD	0.0228 (3.64)	0.0045 (0.47)	0.0075 (0.91)
TimeBudget_CD	0.0008 (2.29)	0.0025 (5.54)	0.0010 (2.18)
Urban_PT	0.9197 (2.67)	1.2217 (2.87)	1.2950 (2.73)
Children_PT	-2.0131 (-9.69)	-1.0410 (-3.77)	-1.6427 (-5.19)
SeasonTkt_PT	1.4801 (5.94)	0.2028 (0.70)	0.3397 (1.03)
TimeBudget_B	0.0028 (9.04)	0.0007 (1.74)	0.0017 (3.41)
Frequency Habit Component			
$\beta_{\text{freq_CD}}$	–	0.0908 (8.06)	–
$\beta_{\text{freq_CP}}$	–	0.0825 (2.49)	–
$\beta_{\text{freq_PT}}$	–	0.1425 (8.80)	–
$\beta_{\text{freq_W}}$	–	0.1153 (8.41)	–
$\beta_{\text{freq_B}}$	–	0.1743 (17.11)	–
Neural Network Habit Component			
α_{CD}	–	–	-1.5097 (-14.11)
α_{CP}	–	–	1.5052 (11.19)
α_{PT}	–	–	-1.0748 (-8.13)
α_{W}	–	–	1.1890 (9.52)
α_{B}	–	–	-0.9683 (-11.24)
NN Internal Params	–	–	277
Goodness-of-Fit Metrics			
Number of Individuals	129	129	129
Estimation Observations (N)	2898	2898	2898
Number of Estimated Parameters (K)	31	36	313
Log-Likelihood at Convergence (LL)	-2061.68	-1589.44	-1133.03
McFadden's Pseudo R^2 (ρ^2)	0.5580	0.6592	0.7571
Bayesian Information Criterion (BIC)	4370.49	3465.86	4761.23

Note: Robust t-statistics in parentheses. Standard deviations of random parameters (σ) are displayed as $\exp(\log \sigma)$. Parameters denoted with '–' are not included in the respective model specification.

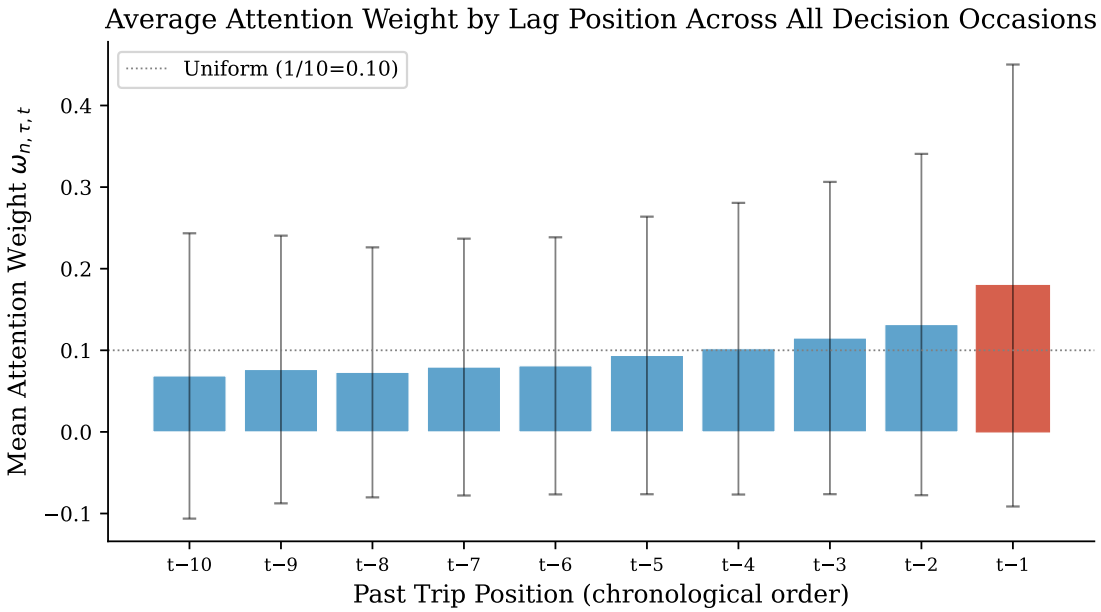


Figure 6: Average attention weight ($\omega_{n,\tau,t}$) by lag position. The dotted line marks the uniform baseline ($1/T = 0.10$) implicitly assumed by traditional frequency counts.

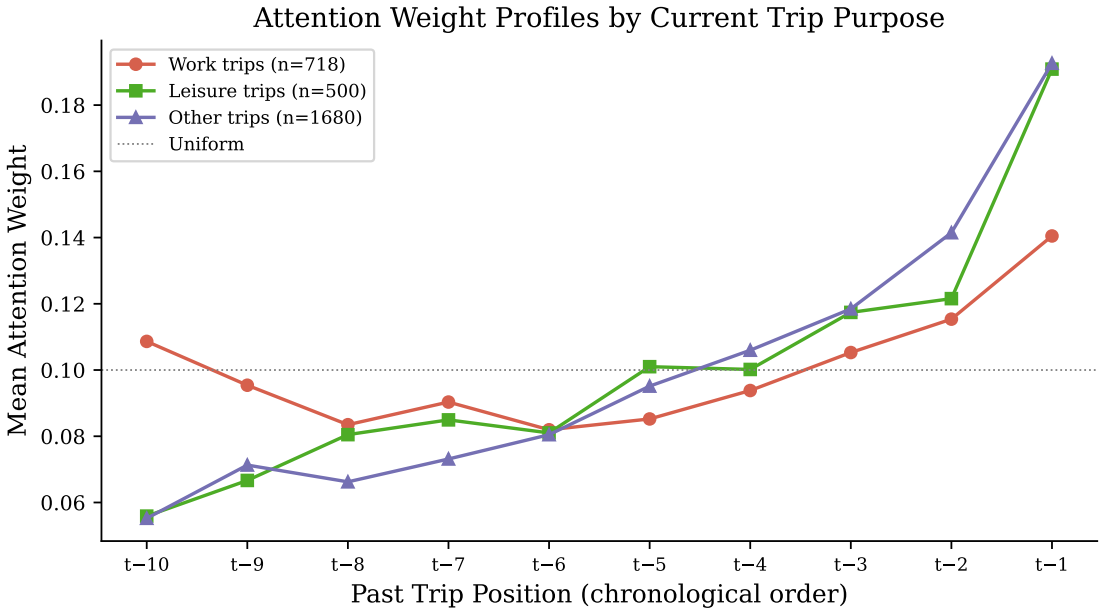


Figure 7: Average attention weight profiles conditioned on the current trip purpose.

lose significance once habit is explicitly modelled, confirming the neural component successfully absorbs these endogenous behavioural effects. Within the random utility framework, the continuous neural habit scores ($\mathbf{H}_{n,t}$) function as a learned index of state dependence. Because this index can take both positive and negative values, similar to traditional inertia variables based on previously experienced utility differences, the estimated parameters (α_i) cannot be interpreted in isolation. Figure 5 displays the distribution of the utility contribution ($\alpha_i H_{in,t}$) across all observations. The habit component predominantly provides positive reinforcement for frequently used modes. As shown in this figure, the car driver mode yields a positive utility contribution in 78% of observations, demonstrating that recent history predominantly reinforces its use.
Figures 6 and 7 reveal the underlying temporal mechanism driving these scores by analysing the network's attention weights ($\omega_{n,\tau,t}$) in Equation 5. Figure 6 shows a clear, data-driven recency gradient: the most recent trip ($t-1$) receives nearly twice the uniform benchmark weight, with influence decaying over older trips. This captures a natural habit decay invisible to frequency-based models, which implicitly assume uniform weighting across all past repetitions. Furthermore, Figure 7 proves this temporal weighting is highly context-dependent. Work trips exhibit a flatter attention profile, drawing heavily on older trips to reflect stable commuting routines. Conversely, leisure trips concentrate sharply on the most recent experiences. From a policy perspective, this distinction is critical: it implies that disrupting habitual commuting may require sustained interventions to shift the accumulated signal, whereas mode-switching for non-routine travel is highly responsive to immediate prior experiences.

4 Conclusion and Future Work

In conclusion, this paper demonstrates that integrating a dynamic neural network habit component into a random utility framework significantly enhances mode choice modelling by capturing context-dependent, temporal habit decay. While these findings offer a behaviourally rich alternative to traditional static frequency measures, several limitations present clear avenues for future econometric and policy-oriented research. First, the empirical application relies on revealed preference data from a stable transport environment. Because transport appraisal requires understanding how habits break, future research should test this neural formulation using panel data encompassing structural shocks or stated preference experiments to observe how the learned habit index adapts when individuals are forced to rethink their choices. Second, to fully operationalise this framework for policy forecasting, it is necessary to develop dynamic simulation procedures that endogenously roll the history window forward under counterfactual scenarios. Finally, while the crossed random error structure accommodates panel correlation, further research using longer-duration panels should investigate the precise interaction between the network's look-back window and unobserved heterogeneity, ensuring that complex state dependence remains robustly disentangled from persistent individual traits. Despite these constraints, this study demonstrates that sequential deep learning provides a principled, practically useful approach to measuring habit, encouraging further integration of data-driven and theory-driven methods in travel behaviour research.

References

Cantillo, V., Ortuzar, J. d. D., & Williams, H. C. (2007). Modeling discrete choices in the presence of inertia and serial correlation. *Transportation Science*, *41*(2), 195–205.

Chen, H., & Cheng, Y. (2023). Travel mode choice prediction using imbalanced machine learning. *IEEE Transactions on Intelligent Transportation Systems*, *24*(4), 3795–3808.

Cherchi, E., & Cirillo, C. (2014). Understanding variability, habit and the effect of long period activity plan in modal choices: a day to day, week to week analysis on panel data. *Transportation*, *41*(6), 1245–1262.

Cherchi, E., Cirillo, C., & Ortúzar, J. de D. (2017). Modelling correlation patterns in mode choice models estimated on multiday travel data. *Transportation Research Part A: Policy and Practice*, *96*, 146–153.

González, R. M., Marrero, Á. S., & Cherchi, E. (2017). Testing for inertia effect when a new tram is implemented. *Transportation Research Part A: Policy and Practice*, *98*, 150–159.

Han, Y., Pereira, F. C., Ben-Akiva, M., & Zegras, C. (2022). A neural-embedded discrete choice model: Learning taste representation with strengthened interpretability. *Transportation Research Part B: Methodological*, *163*, 166–186.

La Paix, L., Oakil, A. T., Hofman, F., & Geurs, K. (2022). The influence of panel effects and inertia on travel cost elasticities for car use and public transport. *Transportation*, *49*(3), 989–1016.

Mo, B., Xu, H., Zhuang, D., Ma, R., Guo, X., & Zhao, J. (2023). Large language models for travel behavior prediction. *arXiv preprint arXiv:2312.00819*.

Omrani, H. (2015). Predicting travel mode of individuals by machine learning. *Transportation research procedia*, *10*, 840–849.

PTV AG, Fell, B., Schönfelder, S., & Axhausen, K. W. (2000, December). *Mobidrive questionnaires* (Vol. 52) [Working Paper]. Zürich, Switzerland: IVT, ETH Zürich. Retrieved from `https://doi.org/10.3929/ethz-a-004233079` (Open access working paper; includes original questionnaires with English translations) doi: 10.3929/ethz-a-004233079

Sifringer, B., Lurkin, V., & Alahi, A. (2020). Enhancing discrete choice models with representation learning. *Transportation Research Part B: Methodological*, *140*, 236–261.

Van Cranenburgh, S., Wang, S., Vij, A., Pereira, F., & Walker, J. (2022). Choice modelling in the age of machine learning-discussion paper. *Journal of choice modelling*, *42*, 100340.

Wang, F., & Ross, C. L. (2018). Machine learning travel mode choices: Comparing the performance of an extreme gradient boosting model with a multinomial logit model. *Transportation Research Record*, *2672*(47), 35–45.

Wang, S., Mo, B., & Zhao, J. (2021). Theory-based residual neural networks: A synergy of discrete choice models and deep neural networks. *Transportation research part B: methodological*, *146*, 333–358.

Wong, M., & Farooq, B. (2021). Reslogit: A residual neural network logit model for data-driven choice modelling. *Transportation Research Part C: Emerging Technologies*, *126*, 103050.

Yáñez, M. F., Cherchi, E., & de Dios Ortúzar, J. (2025). Modelling inertia and shock effects with panel data: An application to mode choice in santiago. *Transportation Research Part A: Policy and Practice*, *199*, 104593.

Yang, Y., Zhang, W., Lin, H., Liu, Y., & Qu, X. (2024). Applying masked language model for transport mode choice behavior prediction. *Transportation Research Part A: Policy and Practice*, *184*, 104074.

Zhang, M., & Xu, Y. (2026). Transmode-llm: Feature-informed natural language modeling with domain-enhanced prompting for travel behavior modeling. *arXiv preprint arXiv:2601.13763*.