Exact Estimation of Activity-Based Dynamic Discrete Choice Models via Reachability-Based Sparse Computation

Azwan Nazamuddin[1], Reem AlOlabi[2], and Makoto Chikaraishi*[1,2]

[1]Graduate School of Innovation and Practice for Smart Society, Hiroshima University, 1-5-1 Kagamiyama, Higashi-Hiroshima, 739-8529, Japan
[2]Graduate School of Advanced Science and Engineering, 1-5-1 Kagamiyama, Higashi-Hiroshima, 739-8529, Japan

*Corresponding author: chikaraishim@hiroshima-u.ac.jp

Short summary

Dynamic Discrete Choice Models (DDCMs) provide a theoretically rigorous, welfare-consistent foundation for activity-based travel demand modeling, but the curse of dimensionality has prevented exact estimation at a practical scale. This study develops a GPU-accelerated computational framework that resolves this barrier through three contributions: reachability-based state space pruning that eliminates the vast majority of physically unreachable states, an activity-sequence partitioned universal graph with constraint-masked batched backward induction that handles population heterogeneity without per-person graph construction, and a time-varying marginal utility specification in which activity timing and sequencing emerge from preference gradients alone. The framework is demonstrated using travel diary data from Higashi-Hiroshima, Japan, with exact NFXP estimation on a single GPU.

Keywords: Activity-Based Models, Dynamic Discrete Choice Models, GPU Acceleration, Nested Fixed Point Estimation, Reachability Analysis, Time-Varying Utility

1 Introduction

Dynamic Discrete Choice Models (DDCMs) provide a theoretically rigorous foundation for activity-based travel demand modeling rooted in microeconomic utility maximization (Västberg et al., 2020), enabling welfare-consistent policy evaluation via the log-sum of value functions (De Jong et al., 2007). However, the curse of dimensionality has severely limited their practical application. Västberg et al. (2020) reported computation times of 4–10 seconds per person for likelihood evaluation, making exact estimation at city scale computationally prohibitive. Consequently, most operational activity-based models have employed rule-based or sequential static approaches, sacrificing explicit intertemporal consistency (Bowman & Ben-Akiva, 2001).

The computational barrier is structural. For a realistic urban setting, the state space encompasses time of day, geographic zone, activity type, duration, transport mode, vehicle availability, and activity history, yielding a theoretical Cartesian product on the order of 10^8 states. Yet the vast majority of these states are physically unreachable from valid initial configurations due to temporal, spatial, and mode-chain constraints (Hägerstraand, 1970; Oyama & Hato, 2019). Naively enumerating the full state space for backward induction is both memory-prohibitive and computationally wasteful, rendering exact nested fixed point (NFXP) estimation (Rust, 1987), which recomputes backward induction at every optimizer iteration, infeasible under conventional approaches.

This study develops a framework that resolves this barrier through three contributions. First, reachability-based state space pruning via GPU-accelerated breadth-first search (Berrendorf et al., 2014) eliminates provably infeasible states prior to dynamic programming, retaining only 1–2% of the theoretical state space while preserving exact optimality. Second, agents sharing the same activity-sequence type (e.g., mandatory work, no mandatory activity) share a single universal graph; individual-level constraints are applied as masks during backward induction rather than requiring separate per-person graphs, substantially reducing computational cost and yielding

O(groups) complexity in the number of population groups. Third, a time-varying marginal utility specification grounded in the utility profile framework of Supernak (1992) and Joh et al. (2003) governs activity timing and sequencing through preference gradients alone: agents depart for work when the work utility gradient exceeds the home reservation floor, and remain at home at night when shopping and leisure utilities fall below it.

Ten behavioral parameters are estimated from a Higashi-Hiroshima travel diary survey ($N = 3{,}331$ person-days) using exact NFXP on a single GPU (NVIDIA RTX 5090).

2 Methodology

DDCM Formulation

The DDCM represents an agent's sequential decision-making over a 24-hour planning horizon. The system state at each step is $s = (\tau, z, a, d, m, v, h)$, where τ is the current time step, $z \in \mathcal{Z}$ is the current zone (144 zones), a is the current activity type (home, work, shopping, leisure, or travel), d is the episode duration since the current activity began, m is the current transport mode (car, walk, bicycle, bus, train), $v \in \{0, 1\}$ is vehicle availability, and h is the mandatory-activity completion bitmask, tracking whether each required activity has been completed. School diary entries are merged into the work activity at the data layer and treated as mandatory activities with the same utility structure. The Cartesian product of these dimensions yields approximately 146 million theoretical states for the Higashi-Hiroshima network.

The dynamic programming solution proceeds via backward induction over the log-sum-exp recursion:

$$V(s) = \log \sum_{a \in \mathcal{A}(s)} \exp Q(s, a;\, \theta) \tag{1}$$

$$Q(s, a;\, \theta) = u(s, a;\, \theta) + V(s') \tag{2}$$

where $u(s, a; \theta)$ is the instantaneous utility under parameters θ and $s' = \text{transition}(s, a)$ is the resulting next state. The terminal condition enforces end-of-day home return:

$$V(s,\, \tau = T) = \begin{cases} 0 & \text{if } s \text{ is a home-terminal state} \\ -\infty & \text{otherwise} \end{cases} \tag{3}$$

Choice probabilities follow the multinomial logit form. The utility $u(s, a; \theta)$ comprises travel disutility (mode-specific constants and time/cost coefficients from an internally estimated mode choice model, unpublished) and activity utility specified as time-varying marginal utility profiles defined in the following section.

Time-Varying Utility Specification

The utility specification follows the temporal utility profile framework of Supernak (1992) and Joh et al. (2003): each activity type has a time-varying marginal utility profile, and an agent pursues an out-of-home activity only when its profile exceeds the home reservation floor. Activity timing, duration, and sequencing emerge organically from these preference gradients.

The model additionally encodes three physical and institutional requirements: (1) mandatory activity completion, enforced via an activity history bitmask in the terminal condition; (2) end-of-day home return, as formalised in the terminal condition of Section 2.1; and (3) vehicle continuity, requiring private vehicles to return home before mode switching.

Home activity accrues a flat marginal utility μ_{home} per minute, serving as the reservation floor that every out-of-home activity must exceed to justify travel.

Work and school follow a piecewise-linear profile around each individual's scheduled hours (t_s, t_e)

from the travel diary:

$$\mu_{\text{work}}(t) = \begin{cases} \delta - \alpha(t_s - t) & t < t_s \quad \text{(earliness penalty)} \\ \delta & t_s \le t \le t_e \quad \text{(on schedule)} \\ \delta - \beta(t - t_e) & t > t_e \quad \text{(lateness penalty)} \end{cases} \tag{4}$$

Near the scheduled start time the earliness penalty fades, the work utility gradient crosses μ_{home}, and the agent departs, with timing emerging from preference gradients rather than any imposed rule.

Shopping and leisure have zone-and-time-varying profiles:

$$\mu(t, z) = \beta_1 \cdot X[z] \cdot P_{\text{open}}[z, t] + \beta_0 \tag{5}$$

where $X[z]$ is zone retail or leisure density and $P_{\text{open}}[z, t]$ is the opening-hour probability at zone z at time t (fit as a Gaussian mixture KDE profile per zone). At night $P_{\text{open}} \to 0$, causing shopping and leisure utilities to fall below μ_{home}, so agents remain at home without any curfew rule.

Travel disutility follows a multinomial logit kernel with mode-specific constants κ_m and time/cost sensitivity coefficients from an internally estimated mode choice model for Higashi-Hiroshima (unpublished):

$$u_{\text{travel}}(m) = \kappa_m + \theta_{\text{travel}} \cdot (\beta_{\text{time}} \cdot t_m + \beta_{\text{cost}} \cdot c_m) \tag{6}$$

where β_{time} and β_{cost} are taken from the internally estimated mode choice model, and θ_{travel} is a scalar multiplier estimated jointly with the activity parameters. This allows the DDCM to freely re-scale travel disutility relative to activity utilities, which is necessary because the utility scales of the mode choice model and the DDCM are not a priori aligned. An activity-switching cost c_{change} is applied at both departure and arrival, imposing a net per-trip overhead of $2c_{\text{change}}$.

This yields ten behavioral parameters in total: δ, α, β, $\beta_{1,\text{shop}}$, $\beta_{0,\text{shop}}$, $\beta_{1,\text{leis}}$, $\beta_{0,\text{leis}}$, c_{change}, μ_{home}, θ_{travel}.

Index	Parameter	Description
0	δ	Work/school on-schedule marginal utility (per min)
1	α	Pre-schedule earliness penalty rate
2	β	Post-schedule lateness penalty rate
3	$\beta_{1,\text{shop}}$	Shopping attractiveness $\times$ P_{open} sensitivity
4	$\beta_{0,\text{shop}}$	Shopping base marginal utility (per min)
5	$\beta_{1,\text{leis}}$	Leisure attractiveness $\times$ P_{open} sensitivity
6	$\beta_{0,\text{leis}}$	Leisure base marginal utility (per min)
7	c_{change}	Activity switching cost (per transition event)
8	μ_{home}	Home flat marginal utility (per min)
9	θ_{travel}	Travel disutility scale (multiplier on MNL time+cost)

Reachability-Based Sparse Computation

A sparse state-transition graph is constructed using level-synchronous breadth-first search (BFS) on GPU (Berrendorf et al., 2014), starting from valid initial states (home at $t = 0$) and expanding only states that satisfy temporal feasibility, space-time prism constraints (Hägerstraand, 1970; Oyama & Hato, 2019), mode-chain rules (e.g., car must return home before switching), and mandatory activity completion requirements. This approach draws on forward reachability analysis from control systems (Bansal et al., 2017): only states reachable from a valid morning configuration are retained, and the rest are discarded before any dynamic programming computation begins.

To avoid per-person graph construction, agents are partitioned by activity-sequence type: the set of mandatory activities required to complete over the day and their structural ordering. Four groups cover the Higashi-Hiroshima population. Agents with *no mandatory activity* face no obligatory out-of-home commitment and may freely allocate the day to shopping, leisure, or home. Agents with *work mandatory* must complete at least one work or school episode, the timing of which is individually specified from the travel diary. The two remaining groups replicate this distinction for agents who additionally have a *child escort* obligation, requiring a separate child drop-off or pick-up trip. Within each group a single universal graph is constructed covering all physically

reachable zone combinations. Individual-level constraints (home zone, mandatory activity timing, opening-hour profiles) are applied as masks during backward induction rather than encoded in the graph structure. Backward induction is then run once per group, batched simultaneously across all relevant home zones in a single GPU pass, giving $O(\text{groups} \times \text{home-zone batches})$ scaling instead of $O(N)$.

The reachable graph is stored in Compressed Sparse Row (CSR) format, with states organized by time step for efficient backward traversal. Value functions are computed via backward induction over this sparse structure using the log-sum-exp operator (Eq. 1), preserving the multivariate extreme value (MEV) choice probability structure and log-sum welfare consistency (De Jong et al., 2007).

Exact Nested Fixed Point Estimation

The computational framework enables exact nested fixed point (NFXP) estimation (Rust, 1987), which embeds full sparse backward induction within the likelihood optimizer. For an observed trajectory $(s_0, a_0, s_1, a_1, \ldots, s_T)$, the log-likelihood is:

$$\log P(\text{path} \mid \theta) = \sum_{t=0}^{T-1} \log P(a_t \mid s_t; \theta) \tag{7}$$

Since $\log P(a_t \mid s_t; \theta) = Q(s_t, a_t; \theta) - \bar{V}(s_t; \theta)$ and value function terms telescope with $\bar{V}(s_T) = 0$, the per-person log-likelihood reduces to (Fosgerau et al., 2013):

$$\ell_n(\theta) = \sum_t u(s_{nt}, a_{nt}; \theta) \; - \; \bar{V}(s_{n,0}; \theta) \tag{8}$$

Taking the derivative, the total gradient separates into observed and expected feature sums (Fosgerau et al., 2013; Ziebart et al., 2008):

$$\frac{\partial \ell(\theta)}{\partial \theta} = \underbrace{\sum_n \sum_t f(s_{nt}, a_{nt})}_{f_{\text{obs}}} \; - \; \underbrace{\sum_n \frac{\partial \bar{V}(s_{n,0})}{\partial \theta}}_{f_{\text{exp}}} \tag{9}$$

where $f(s, a) = \partial u(s, a)/\partial \theta$ is the sparse feature vector (linear-in-parameters utility). The f_{obs} term depends only on observed state-action pairs and does not require model-based expectations or backward recursion; f_{exp} requires $\partial \bar{V}(s_{n,0})/\partial \theta$ per person, computed via a second backward induction pass over the same Bellman Directed Acyclic Graph (DAG) (Baydin et al., 2018):

$$\frac{\partial \bar{V}(s)}{\partial \theta} = \sum_{a \in \mathcal{A}(s)} P(a \mid s; \theta) \left[f(s, a) + \frac{\partial \bar{V}(s')}{\partial \theta} \right] \tag{10}$$

with base case $\partial \bar{V}(s_T)/\partial \theta = 0$. This is the *observed minus expected features* gradient applied to finite-horizon DAGs (Fosgerau et al., 2013; Ziebart et al., 2008).

Persons sharing the same activity-sequence group share a single backward induction computation, enabling amortization across the sample. Graph construction runs once and is retained; only backward induction and likelihood evaluation repeat within each optimizer step.

Optimization uses BFGS with the analytical gradient derived above. GPU-based parallel computation introduces small but irreducible numerical noise in the log-likelihood evaluation (Shanmugavelu et al., 2024, 2025); approximating the gradient numerically would amplify this noise to the point where the optimizer cannot distinguish a meaningful search direction from numerical error. The analytical gradient, which is computed exactly from the same backward induction pass, avoids this amplification entirely and is therefore necessary for stable convergence on GPU. Two numerical stabilization measures are applied to make BFGS tractable: parameters are rescaled so that all ten are of comparable magnitude at the starting point, and the gradient is normalized by its initial scale so that early iterations do not produce excessively large steps on the high-dimensional state-action graph. As a byproduct, the per-person gradient vectors produced during estimation yield BHHH standard errors (Fosgerau et al., 2013) at no additional computational cost.

Case Study: Higashi-Hiroshima

The framework is applied to Higashi-Hiroshima City, Hiroshima Prefecture, Japan, a mid-size city with JR rail, bus networks, and high car and bicycle ownership, comprising 144 analysis zones and 5 transport modes (car, walk, bicycle, bus, train). The data source is a one-day household travel diary survey. After filtering for valid observed paths and mandatory activity constraints, the estimation sample comprises $N = 316$ person-days. All computation uses a single NVIDIA RTX 5090 GPU (32 GB VRAM).

3 Results and Discussion

Computational Performance

Table 1 reports measured performance on the Higashi-Hiroshima 144-zone real network. Reachability-based pruning retains only 1–2% of the theoretical state space across all four activity-sequence groups while preserving exact optimality. The universal graph architecture eliminates per-person graph construction; the entire population is covered by four graphs with batched backward induction across home zones. Graph construction runs once and is reused across all NFXP optimizer iterations; only backward induction and likelihood evaluation repeat at each step.

Table 1: Computational Performance of the GPU-Accelerated Pipeline (RTX 5090, 32 GB)

Activity-sequence group	**Reachable states**	**Theoretical**	**Reduction**
No mandatory activity	5.84 M	146 M	96.0%
Work mandatory	2.39 M	146 M	98.4%
Phase	**Time (RTX 5090)**		
Graph construction (per group)	16–34 s		
Batched backward induction (per group)	0.7–3.8 s		
Forward simulation (1,000 agents)	< 30 s		
End-to-end, all groups (graph $\rightarrow$ simulate)	$\sim$2.5–3 min		
Per person per evaluation	$\ll 0.1$ s (cf. 4–10 s[a])		
Memory (dense Q-table / sparse CSR)	$\sim$6.7 TB vs. 5–10 GB	($\sim$1,000$\times$ reduction)	

[a] Västberg et al. (2020), CPU-based per-person figure.

Behavioral Validation

Behavioral validation uses a 3-zone synthetic network (Zone 1: home; Zone 2: work; Zone 3: shopping/leisure) with manually specified travel times and Gaussian-mixture opening-hour profiles, simulating 300 agents per scenario. Figure 1 shows the marginal utility profiles and the resulting simulated activity fractions: work timing emerges from the piecewise-linear profile crossing μ_{home}, shopping follows the P_{open} Gaussian mixture profile, and leisure activates only above its utility threshold. The toy case confirms that the marginal utility specification produces internally consistent activity patterns from utility gradients alone. Table 2 summarises the baseline outcomes for each activity type.

Estimated Parameters

Table 3 presents the parameter estimates from exact NFXP on 316 workers drawn by simple random sampling from the Higashi-Hiroshima travel diary survey, covering 17,933 observed activity-travel steps. The travel disutility is taken from a separately estimated mode choice model and enters the dynamic model with a single scaling parameter, ensuring consistency of relative mode preferences while allowing for rescaling in the dynamic context.

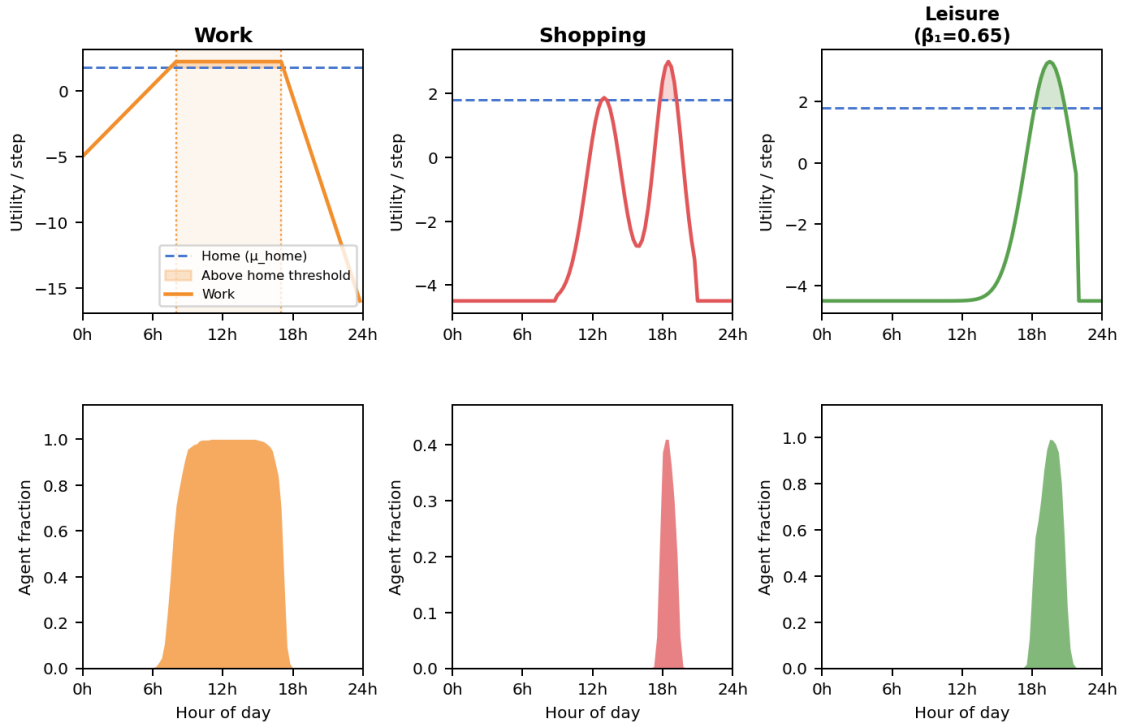


Figure 1: Marginal utility profiles (top row) and simulated activity fraction (bottom row) for three activity types under the marginal utility specification. Work timing emerges from the piecewise-linear profile crossing the home reservation floor μ_{home}; shopping follows the P_{open} Gaussian mixture profile; leisure activates above its utility threshold ($\beta_{1,\text{leis}} = 0.65$). Conceptually grounded in Supernak (1992).

Table 2: Behavioral Validation: Toy Case Baseline Outcomes (marginal utility specification, 3-zone synthetic network)

Activity	**Timing source**	**Baseline outcome**
Work	Piecewise marginal utility around t_s, t_e	Workers at work 8–17h; utility-emergent, no rule
Shopping	P_{open} Gaussian mixture profile	After-work shopping; suppressed at cost $\times 1.58$
Leisure	P_{open}, activation threshold $\beta_1 \approx 0.525$	Inactive at baseline; activates above threshold
Home	μ_{home} reservation floor	Dominant for non-workers; welfare floor for all

The results show the travel scale $\hat{\theta}_{\text{travel}} = 1.014$, which is nearly identical to the internally estimated mode-choice baseline. The shopping and leisure density weights ($\hat{\beta}_{1,\text{shop}} = 0.032$, $\hat{\beta}_{1,\text{leis}} = 0.002$) are positive, consistent with out-of-home activities becoming more attractive in denser, open zones; the negative base utilities ($\hat{\beta}_{0,\cdot}$) imply that zone density and opening hours must be sufficiently high before either activity exceeds the home floor.

The estimates in Table 3 reflect a local optimum rather than the global MLE. To understand why, it is useful to consider how the switching cost parameter interacts with the rest of the model. The cost c_{change} is incurred at both departure and arrival, so each activity transition carries a net overhead of $2c_{\text{change}}$. When this cost is large in magnitude, the model can reproduce almost any observed schedule through inertia rather than through calibrated activity preferences: agents rarely change activities not because they lack attractive alternatives, but because the transition overhead dominates any utility gain. In this regime the individual activity parameters become weakly identified, since predicted behavior changes little when their values are varied. The likelihood surface develops a flat region along the activity-utility dimensions, and the optimizer is unable to progress further. This is reflected in the gradient norm stabilising at $\|\nabla \ell\|_\infty = 0.616$ from iteration 7 onward, well above the convergence threshold of 0.001, and in several parameters taking atypical

Table 3: NFXP Parameter Estimates (marginal utility specification, $N = 316$, $T = 17{,}933$ observed steps)

Parameter	Description	Estimate
Work/school schedule utility		
δ	On-schedule marginal utility (per min)	-0.0235
α	Pre-schedule earliness penalty rate	0.000855
β	Post-schedule lateness penalty rate	8.0×10^{-6}
Shopping		
$\beta_{1,\text{shop}}$	Retail density $\times$ P_{open} sensitivity	0.0317
$\beta_{0,\text{shop}}$	Base marginal utility (per min)	-0.0428
Leisure		
$\beta_{1,\text{leis}}$	Leisure density $\times$ P_{open} sensitivity	0.00188
$\beta_{0,\text{leis}}$	Base marginal utility (per min)	-0.0411
Activity switching and home		
c_{change}	Switching cost (per departure/arrival)	-3.862
μ_{home}	Home flat marginal utility (per min)	-0.0254
Travel		
θ_{travel}	Travel disutility scale	1.014
Log-likelihood $\ell(\hat{\theta})$		$-3{,}357.2$
BFGS iterations (best checkpoint)		7 (of 33 total)

Mode constants and time/cost coefficients from an internally estimated mode choice model (unpublished); only θ_{travel} is estimated jointly. Gradient norm $\|\nabla\ell\|_\infty = 0.616$ at best checkpoint (convergence threshold: 0.001); estimates reflect a local optimum, as discussed in the text.

signs: $\hat{\delta} < 0$ implies negative on-schedule work utility, $\hat{\mu}_{\text{home}} < 0$ means the home floor acts as a mild penalty rather than a positive reservation level, and the near-zero lateness penalty $\hat{\beta}$ is effectively unidentified.

The computational cost of each BFGS iteration compounded this difficulty. Backward induction must be run sequentially across all scheduling groups, each comprising roughly 5.9 million reachable states and 1.86 billion state-action edges and requiring approximately 7 GB of GPU memory per pass. At around 22 minutes per iteration, the available overnight compute window before submission supported at most 18 to 20 iterations per run. Two full runs totalling 33 iterations were completed, exploring two initialization strategies: a warm start from a prior checkpoint, and a partial reset of the switching cost. Neither reached the global MLE, and the best result across all 33 iterations remained at iteration 7 of the first run. Multi-start exploration from fresh, economically grounded initial values requires at least three to five independent runs, which the submission timeline did not allow.

The results reported here therefore represent a partial convergence outcome. The NFXP pipeline is fully operational and produces stable log-likelihood improvement in early iterations; the remaining gap concerns optimizer initialization and iteration budget rather than any fundamental methodological limitation. Two directions would address this in future work. Constraining the switching cost to a moderate range from the outset prevents the degenerate basin without additional computation. Coarser timing discretization would reduce per-iteration wall time to roughly 7 to 8 minutes, making 50 to 100 iterations feasible within a single overnight window and allowing multi-start exploration within a pre-submission compute budget.

4 Conclusions

This study demonstrates that exact DDCM estimation is computationally feasible at the scale of a city-level travel diary survey. Three contributions make this practical. First, reachability-based pruning retains only 1–2% of the theoretical state space while preserving exact optimality, making backward induction tractable on a single GPU. Second, an activity-sequence partitioned universal graph with constraint-masked batched backward induction eliminates per-person graph

construction, covering the full Higashi-Hiroshima population with four graphs and achieving end-to-end simulation in under three minutes. Third, a time-varying marginal utility specification grounded in the utility profile framework of Supernak (1992) and Joh et al. (2003) governs activity timing, sequencing, and trip-making through preference gradients alone, with the home reservation floor μ_{home} as the behavioral anchor.

Toy-case validation confirms that all three activity types exhibit utility-emergent timing consistent with Supernak (1992). NFXP estimation from real Higashi-Hiroshima survey data yields a partial convergence result, demonstrating that the pipeline is fully operational and produces stable log-likelihood improvement in early iterations; remaining challenges concern optimizer initialization and iteration budget, as discussed in Section 3.3. Future work will pursue fully converged estimates and extend the framework to welfare-consistent policy evaluation via log-sum consumer surplus (De Jong et al., 2007).

Acknowledgements

The authors would like to acknowledge that a part of this research was partially supported by JST FOREST (JPMJFR231G), and Council for Science, Technology and Innovation (CSTI), Cross-ministerial Strategic Innovation Promotion Program (SIP), the 3rd period of SIP "Smart Mobility Platform", Grant Number JPNP23023 (Funding agency: NEDO).

References

Bansal, S., Chen, M., Herbert, S., & Tomlin, C. J. (2017, September). Hamilton-Jacobi Reachability: A Brief Overview and Recent Advances [arXiv:1709.07523 [eess]]. https://doi.org/10.48550/arXiv.1709.07523

Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018, February). Automatic differentiation in machine learning: A survey [arXiv:1502.05767 [cs]]. https://doi.org/10.48550/arXiv.1502.05767

Berrendorf, R., Makulla, M., & Augustin, S. (2014). Level-Synchronous Parallel Breadth-First Search Algorithms For Multicore and Multiprocessor Systems.

Bowman, J., & Ben-Akiva, M. (2001). Activity-based disaggregate travel demand model system with activity schedules. *Transportation Research Part A: Policy and Practice*, *35*(1), 1–28. https://doi.org/10.1016/S0965-8564(99)00043-9

De Jong, G., Daly, A., Pieters, M., & Van Der Hoorn, T. (2007). The logsum as an evaluation measure: Review of the literature and new results. *Transportation Research Part A: Policy and Practice*, *41*(9), 874–889. https://doi.org/10.1016/j.tra.2006.10.002

Fosgerau, M., Frejinger, E., & Karlstrom, A. (2013). A link based network route choice model with unrestricted choice set. *Transportation Research Part B: Methodological*, *56*, 70–80. https://doi.org/10.1016/j.trb.2013.07.012

Hägerstraand, T. (1970). WHAT ABOUT PEOPLE IN REGIONAL SCIENCE? *Papers in Regional Science*, *24*(1), 7–21. https://doi.org/10.1111/j.1435-5597.1970.tb01464.x

Joh, C.-H., Arentze, T. A., & Timmermans, H. J. P. (2003). Understanding Activity Scheduling and Rescheduling Behaviour: Theory and Numerical Illustration [Series Title: The GeoJournal Library]. In M. Barlow, B. Boots, A. Okabe, & R. Thomas (Eds.), *Modelling Geographical Systems* (pp. 73–95, Vol. 70). Springer Netherlands. https://doi.org/10.1007/978-94-017-2296-4_5

Oyama, Y., & Hato, E. (2019). Prism-based path set restriction for solving Markovian traffic assignment problem. *Transportation Research Part B: Methodological*, *122*, 528–546. https://doi.org/10.1016/j.trb.2019.02.002

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, *55*(5), 999. https://doi.org/10.2307/1911259

Shanmugavelu, S., Taillefumier, M., Culver, C., Ganesh, V., Hernandez, O., & Sedova, A. (2025, August). Robustness of deep learning classification to adversarial input on GPUs: Asynchronous parallel accumulation is a source of vulnerability [arXiv:2503.17173 [cs]]. https://doi.org/10.48550/arXiv.2503.17173

Shanmugavelu, S., Taillefumier, M., Culver, C., Hernandez, O., Coletti, M., & Sedova, A. (2024, October). Impacts of floating-point non-associativity on reproducibility for HPC and deep learning applications [arXiv:2408.05148 [cs]; Published: IEEE SC'24 (DOI: 10.1109/SC41406.2024.00059)]. https://doi.org/10.48550/arXiv.2408.05148

Supernak, J. (1992). Temporal utility profiles of activities and travel: Uncertainty and decision making. *Transportation Research Part B: Methodological*, *26*(1), 60–76. https://doi.org/10.1016/0191-2615(92)90020-W

Västberg, O. B., Karlström, A., Jonsson, D., & Sundberg, M. (2020). A Dynamic Discrete Choice Activity-Based Travel Demand Model. *Transportation Science*, *54*(1), 21–41. https://doi.org/10.1287/trsc.2019.0898

Ziebart, B. D., Maas, A., Bagnell, J. A., & Dey, A. K. (2008). Maximum Entropy Inverse Reinforcement Learning. *AAAI Press*.