

Human-Flow Digital Twin for Predicting the Effects of Mobility Introduction on Visitor Circulation

Chiharu Shima^{ID*}, Haruki Yonekura^{ID†‡}, Fukuharu Tanaka^{ID†‡}, Tatsuya Amano^{ID†‡}, Hirozumi Yamaguchi^{ID†‡}

^{*}bitA Inc., Shinagawa-ku, Tokyo, Japan

[†]The University of Osaka, Suita, Osaka, Japan

[‡]RIKEN Center for Computational Science, Kobe, Hyogo, Japan

{c-shima, h-yonekura, f-tanaka, t-amano, h-yamagu}@ist.osaka-u.ac.jp

Abstract—We propose a framework for predicting the effects of mobility introduction measures using a human-flow digital twin. This digital twin incorporates a multi-agent simulator that can represent how visitors choose destinations depending on factors such as their current location and the attractiveness of spots. We extract data on how visitors selected destinations with respect to measured pre-intervention human-flow data, inter-spot distances, spot attractiveness, and travel volumes, and use these data to train each agent’s decision model of this simulator. The trained decision model is a function that takes a visitor’s current state and surrounding environmental information as input and outputs which spot the visitor will move toward next. By expressing mobility introduction measures as changes to inter-point distances or to spot attractiveness, the framework can reproduce human flows with mobility introduction in the multi-agent simulator and thereby quantify effects such as changes in visitor counts and circulation. We evaluated the proposed method using human-flow data measured with and without introducing mobility within Wakayama Castle Park in Japan. When reproducing flows with mobility introduction using a multi-layer perceptron decision model, the cosine similarity of the spatial population distribution exceeded 0.7, confirming that the approach can replicate the flow changes caused by the mobility introduction.

Index Terms—digital twin, human-flow, agent-based simulation, decision-making model, tourism management.

I. INTRODUCTION

Tourism promotion is widely recognized as an effective means of regional revitalization. Local visitor destinations often possess valuable cultural and historical assets, such as heritage buildings, museums, and scenic sites, that can attract visitors. However, limited connectivity and the relative difficulty of moving between attractions frequently impede visitor circulation, reducing opportunities for visitors to explore multiple sites and undermining the overall attractiveness and economic potential of these areas. To mitigate this problem, various new mobility options have been proposed and piloted to facilitate intra-site movement. Personal mobility devices, such as e-scooters and e-bikes, and green slow-mobility services, such as low-speed electric shuttle carts and golf-cart-like vehicles, are among the approaches expected to enhance visitor circulation by reducing travel friction and making inter-site movement more convenient [1], [2].

This paper has been accepted for presentation at the 27th IEEE International Conference on Mobile Data Management(MDM2026).




Fig. 1: Predicting the Effect of Mobility Introduction via Digital Twin Simulation.

When planning mobility introductions at visitor sites, decision-makers must determine operational parameters such as routing, the spatial placement of boarding and alighting points, service frequency, and fleet size that will most effectively increase visitor circulation. Conducting field trials for every candidate design is, however, prohibitively costly and time-consuming. Consequently, there is a practical need for predictive tools that can evaluate alternative deployment strategies via simulation before physical implementation, thereby guiding planners toward effective and cost-efficient solutions. Although prior work has advanced destination-choice modeling and microscopic pedestrian simulation [3], [4], few studies explicitly combine learned destination-choice models with region-scale agent-based simulators to evaluate the impacts of introducing new mobility options.

In practice, introducing low-speed mobility services into visitor sites is not merely a matter of adding vehicles. Historic parks and cultural districts often contain narrow walkable spaces and mixed-use circulation, so new mobility modes can intensify competition for limited urban space and create site-specific operational and safety challenges. Prior work on micro-mobility integration has emphasized that municipalities are often unprepared for such interactions and that data-driven evaluation is important for devising acceptable regulations and countermeasures [5]. These concerns are particularly salient in tourism settings, where boarding and alighting locations, route design, and pedestrian congestion must be coordinated without undermining visitor comfort or accessibility.

From a planning perspective, the problem is inherently counterfactual: the same site cannot be observed under mul-

multiple deployment designs, and reliable intervention-specific data are rarely available before implementation. Recent studies have begun to explore counterfactual prediction for tourism planning under privacy-sensitive observations [6] and to assess the spatial effects of green mobility in urban parks from limited pre-deployment information [7]. More broadly, digital twins are attracting increasing attention in transport planning, but their effective application to complex planning processes remains challenging, especially when behavioral adaptation and scenario-based intervention assessment must be integrated within a usable planning workflow [8]. These considerations motivate a framework that learns visitor responses from normal-state observations and evaluates mobility introduction as a controlled environmental modification rather than as a separately observed intervention dataset. In the present implementation, offline trajectory reconstruction and simulation initialization rely on normal-state wide-area location-point sequences that are privacy-safe and anonymized/pseudonymized, but still retain tracking IDs and timestamps, together with site-level spot counts and environmental information. Following Uegaki et al. [9], these location-point sequences are used to estimate inter-PoI movement priors and relative Origin-Destination (OD) composition for offline initialization, while intervention-specific retraining is avoided at the online prediction stage.

In this paper, we propose a platform that predicts and visualizes the effects of *mobility introduction* (i.e., the deployment of mobility services, such as e-scooters and low-speed electric shuttle carts on designated routes within a visitor site, enabling visitors to move between points of interest) on visitor circulation by means of a human-flow digital twin shown in Fig. 1. The core of the platform is an agent-based simulation (ABS) that explicitly models the link between individual decision-making and environmental factors. Each visitor is represented as an autonomous agent whose destination-choice behavior is governed by a learned decision model. Given an agent’s internal state (current location, previously visited points of interest (PoIs), and cumulative walking distance) and environmental inputs (pairwise distances between PoIs, PoI attractiveness, and time of day), the model computes a probability distribution over the agent’s candidate next destinations. We treat mobility introduction as modifications to the environmental factors that feed into the decision models; by modeling mobility indirectly through such environmental modifications, our approach obviates the need for large, intervention-specific mobility datasets. By applying the learned decision models within the ABS under such modified environmental conditions, the platform can simulate how the mobility introduction—on selected routes and with specified service parameters—affects visitors’ route choices and, in aggregate, the spatio-temporal distribution of visitors. For the simulation platform, we integrated Simulation of Urban MObility (SUMO) [10], JuPedSim [11], and Traffic Control Interface (TraCI) API [12] to capture both high-level decision making and fine-grained pedestrian dynamics.

We validated the proposed framework using human-flow data collected within Wakayama Castle Park in Japan. The

park area was partitioned into eight areas, and we evaluated the platform’s ability to reproduce the temporal evolution of per-area population distributions. The results indicate that the Multi-Layer Perceptron (MLP) based decision model can reproduce the temporal population distribution under mobility introduction with cosine similarity exceeding 0.7. In addition, the MLP model reproduced the mobility-induced change in population distribution with a cosine similarity of 0.664.

II. RELATED WORKS

A. Destination-Choice Modeling

Research on destination-choice modeling seeks to explain and predict how individuals select destinations by relating observable environmental factors and latent psychological drivers to choice behavior. Robin et al. [13] propose discrete-choice formulations that estimate decision probabilities from observational data to model rational destination-selection processes. Gao et al. [14] construct a tourism-related knowledge graph and introduce a Knowledge-Graph-aware Disentangled Auto-Encoder (KGDAE) to model tourist choices based on geographic and functional factors. Dai et al. [15] define latent motivational constructs such as “travel inspiration” and use these psychological factors to model traveler decision processes. McCabe et al. [16] develop a general model grounded in dual-system theory that emphasizes the role of psychological context and the interplay of personal and contextual factors in shaping choice strategies. King et al. [17] conduct virtual experiments to investigate whether destination-choice behavior depends on building layout and planned or imposed destination schedules. Jin et al. [18] apply machine-learning models and use SHapley Additive exPlanations (SHAP) to analyze and interpret pedestrian route-choice mechanisms. While these works advance understanding of microscopic choice mechanisms and provide reliable building blocks for demand modeling, they are typically developed and evaluated in settings focused on modeling choice behavior. In particular, the literature rarely demonstrates end-to-end propagation of learned destination-choice models into physics-aware, region-scale pedestrian simulations for the purpose of evaluating mobility-mode introductions.

B. Pedestrian Microsimulation

Microscopic pedestrian simulation research focuses on detailed local movement rules and interaction models that reproduce short-term crowd dynamics and emergent collective behaviors. Helbing et al. [19] develop the social force family of models, which represent pedestrian motion through attractive and repulsive interactions and explain emergent patterns such as lane formation and bottleneck oscillations. Treuille et al. [20] propose a continuum-based, real-time crowd simulation method that couples global navigation with local collision avoidance via dynamic potential fields. Armel et al. [11] present JuPedSim, an open-source framework for simulating and analyzing microscopic pedestrian dynamics. Zhang et al. [21] introduce a speed-based pedestrian model that integrates

walking comfort and directional preferences into a collision-free framework. Dang et al. [22] propose an adaptive scheme that switches local motion models according to nearby density to better reproduce individual trajectories. Such microsimulators are widely used for safety and operational analysis; however, they frequently assume exogenous routing or demand inputs and are not always designed to incorporate empirically-learned, context-sensitive destination-choice behavior directly.

C. Agent-Based Simulation (ABS)

ABS research aims to bridge individual decision processes and population-level movement by embedding explicit choice models within simulated agents. Huang et al. [23] extend floor-field cellular automata to model decision making under risk and uncertainty in evacuation scenarios. Horni et al. [24] augment MATSim with a high-resolution destination-choice formulation that embeds random-utility components to enable realistic location choice in large-scale agent simulations. Wörle et al. [25] model intermodal travel behavior and its influence on destination and mode choice within an agent-based travel-demand framework. Kazyeva et al. [3] develop a region-scale agent-based simulation that incorporates walkability considerations to better reflect heterogeneous behavior of tourists and residents. Yang et al. [26] propose a co-simulation system that couples MATSim with external fleet and mode simulators to propagate destination-based demand through multimodal simulations. Li et al. [27] describe procedures to generate and validate large-scale agent-based transport models against observed mode shares and travel-time distributions. Nevertheless, many ABS implementations focus either on macroscopic demand propagation or on transport-mode interactions, and they are not routinely configured to evaluate the introduction of novel, low-speed personal mobility services in a way that jointly preserves micro-level physical interactions and macro-level destination dynamics.

D. Contribution of This Study

This paper addresses two practical gaps in the related works: the relatively limited attention to modelling the mobility introduction and their impact on visitor circulation, and the lack of a unified, implementable framework that integrates learned destination-choice models with physics-aware microsimulation and end-to-end visualization for site-scale intervention assessment. The main contributions of this paper are summarized as follows.

- We present a human-flow digital twin that embeds learned destination-choice models within an ABS, enabling the reproduction and the visualization of circulation with mobility introduction.
- We model mobility introduction as changes in environmental features and inject them into learned decision models to predict individual choices. This reduces the need for intervention-specific data, although offline initialization still requires normal-state anonymized/pseudonymized wide-area location se-

quences with tracking IDs and timestamps, along with spot counts, to estimate movement priors.

- We deliver a practical simulation platform that couples SUMO, destination-choice models and a Social Force model, thereby preserving both macro-level routing and fine-grained local crowd interactions.
- We provide empirical validation using Wakayama Castle Park data collected through a field experiment covering both the park and adjacent public roads: the MLP model reproduces per-area temporal population trajectories with cosine similarity values exceeding 0.7 and reproduces mobility-induced distribution changes with a cosine similarity of 0.664.

III. PROPOSED METHOD


A. Overview of the Proposed Method

We propose a platform that predicts and visualizes how visitor behavior and aggregate flows change after the mobility introduction. Figure 2 provides an overview of the pipeline. In this study, we focus on low-speed, single-occupant personal mobility devices (e.g., electric kick scooters and small electric carts) intended for intra-site circulation. These devices are modeled as operating on a set of predefined mobility links that connect fixed segment endpoints; riders may board or alight only at these designated endpoints. The platform takes as input (i) normal-state anonymized/pseudonymized wide-area location-point sequences with tracking IDs and timestamps, (ii) time-stamped pedestrian traffic volumes measured at selected counting spots (counters), and (iii) environmental information such as the walkable network, Points of Interest (PoIs), and indicators of mobility introduction. The output is a simulation scenario and a synthesized set of agents' trajectories whose aggregate characteristics are consistent with the measured spot volumes. The resulting trajectories are mapped onto a 3D map using Unity and Cesium to support geographic inspection of simulation outcomes. Agents are instantiated to match the reconstructed departure events. At each decision step, an agent's next-destination is updated by a destination-choice model, and the agent is then advanced in the simulator accordingly, yielding an agents' trajectory. Each decision model maps an agent's internal state (e.g., current location, time of day, and accumulated travel distance) and environmental inputs (e.g., pairwise PoI distances, PoI attractiveness, and movement costs) to a probability distribution over candidate next destinations. By switching the environmental inputs to represent mobility introduction (e.g., reduced generalized travel cost or increased attractiveness at mobility-enabled locations), the platform reproduces how agents' choices adapt and how these adaptations aggregate into population-level flow changes.

B. Offline Phase: Data Preparation from Wide-Area Location Data and Spot Count

1) *Observation Model for Spot Counts and Wide-Area Location Sequences:* Following the pedestrian flow generation

Offline Phase: Model Training(Usual Situation).



Online Phase: Prediction on new environment with Mobility.




Fig. 2: Overview of the Proposed Platform.

method of Uegaki et al. [9], we introduce a structured movement prior estimated from wide-area location point sequences. Wide-area location data refers to privacy-preserving location-point sequences collected from mobile devices or location-based services across a broad spatial region. Such data are typically anonymized, coarsened in space or time, and sometimes perturbed or filtered to protect user privacy. Consequently, they mainly describe aggregate inter-area movement tendencies rather than exact agents' trajectories.

The proposed pipeline combines normal-state wide-area location-point sequences with pedestrian traffic volumes measured at a subset of areas within a region of interest. Let $\mathcal{A}$ denote the set of all areas covering the region of interest, and let $\mathcal{A}^{\text{obs}} \subseteq \mathcal{A}$ denote the subset of areas where spot counts are available. The location-point sequences are map-matched to the area network, yielding origin–destination observations over all ordered area pairs under normal conditions, whereas the spot counts provide accurate site-specific constraints on local flow intensity at $\mathcal{A}^{\text{obs}}$.

Measurements are aggregated into discrete time slots; let $\mathcal{T} = \{1, 2, \dots, T\}$ be the set of slots, let Δ be the slot duration (e.g., $\Delta = 10$ minutes), and denote the interval of slot $t \in \mathcal{T}$ by $I_t = [(t-1)\Delta, t\Delta)$.

For each observed area $a \in \mathcal{A}^{\text{obs}}$ and slot $t \in \mathcal{T}$, we record the number of pedestrians passing the counting spot located in area a during I_t , denoted by $L_{a,t}$. Formally, letting $e_{i,a}(\tau)$ indicate that pedestrian i passes the counting spot in area a at time τ , we define

$$L_{a,t} = \sum_i \mathbf{1}[\exists \tau \in I_t \text{ s.t. } e_{i,a}(\tau) = 1], \quad (1)$$

where $\mathbf{1}[\cdot]$ is the indicator function. Eq. 1 is a conceptual definition; in practice, we observe only the aggregated counts $L_{a,t}$ without identity-level linkage. This yields a time-stamped

dataset $\{L_{a,t}\}_{a \in \mathcal{A}^{\text{obs}}, t \in \mathcal{T}}$, which is used in the subsequent flow reconstruction module.

In addition, the wide-area location-point sequences are map-matched to obtain OD observations at the area level. Let (o_i, d_i) denote the origin and destination areas of movement sample i after map matching, and let $v_i \in V$ denote the speed category assigned to that trajectory. We define the observed OD counts as

$$X_{m,n,t,v} = \sum_i \mathbf{1}[o_i = m, d_i = n, \tau_i \in I_t, v_i = v], \quad (2)$$

where τ_i is the departure time of movement sample i . This yields a time-stamped OD dataset $\{X_{m,n,t,v}\}_{m,n \in \mathcal{A}, t \in \mathcal{T}, v \in V}$. The same map matched trajectory samples are also used before aggregation to estimate the continuous departure time and travel duration distributions described in Sec. III-B2. In the reconstruction stage, the aggregated counts provide the empirical relative OD composition. Specifically, for each ordered area pair (m, n) , we define

$$c_{m,n} = \sum_{t \in \mathcal{T}} \sum_{v \in V} X_{m,n,t,v}, \quad (3)$$

which represents the empirical number of observed transitions from area m to area n in the auxiliary location dataset.

2) *Flow Reconstruction and Calibration:* Given the spot-volume observations $\{L_{a,t}\}$ defined in Sec. III-B1, together with auxiliary normal-state mobility observations used to estimate inter-area priors, we reconstruct a statistical departure process and synthesize trajectories consistent with these aggregate measurements. We use a prior-demand generation and scaling procedure. Following prior count-based pedestrian-flow reconstruction methods [9], we estimate plausible pedestrian departures while acknowledging the non-uniqueness of the reconstruction.

Because the spot counts provide only partial aggregate observations of the full pedestrian flow, they are insufficient to determine all area-to-area departures by themselves. We therefore first construct a structured prior demand scenario from the wide-area location data and then calibrate its absolute scale using the observed spot counts.

For each ordered area pair $(m, n) \in \mathcal{A} \times \mathcal{A}$, we model the joint distribution of departure time of day and travel duration using a 2D Gaussian mixture model (GMM):

$$p(\mathbf{z} \mid m, n) = \sum_{k=1}^K \pi_k^{(m,n)} \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_k^{(m,n)}, \boldsymbol{\Sigma}_k^{(m,n)}) \\ := G_{m,n}(\mathbf{z}), \quad \mathbf{z} = [t^{\text{dep}} \ d^{\text{trav}}]^\top, \quad (4)$$

where t^{dep} denotes the departure time of day and d^{trav} denotes the travel duration in minutes. Here, K is the number of Gaussian components, and $\pi_k^{(m,n)}$ is the mixture weight of component k for movements from area m to area n , satisfying $\pi_k^{(m,n)} \geq 0$ and $\sum_{k=1}^K \pi_k^{(m,n)} = 1$. The term $\mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_k^{(m,n)}, \boldsymbol{\Sigma}_k^{(m,n)})$ denotes a bivariate Gaussian density over the two-dimensional vector $\mathbf{z}$. The mean vector $\boldsymbol{\mu}_k^{(m,n)} \in \mathbb{R}^2$ represents the typical departure time and travel duration of component k , while the covariance matrix $\boldsymbol{\Sigma}_k^{(m,n)} \in \mathbb{R}^{2 \times 2}$ captures their variances and correlation. Thus, each component can represent a characteristic movement pattern, such as short-duration trips during one period of the day or longer-duration trips during another period.

In practice, the set of priors $\{G_{m,n}\}$ is estimated offline from the unaggregated map matched trajectory samples used to construct Eq. (2). The aggregated OD counts $\{X_{m,n,t,v}\}$ provide the empirical transition counts $\{c_{m,n}\}$ in Eq. (3), which preserve the relative OD composition during sampling. The estimated distributions are then treated as fixed priors during calibration.

For each ordered area pair (m, n) , we draw $c_{m,n}$ independent samples $(t_i^{\text{dep}}, d_i^{\text{trav}}) \sim G_{m,n}$ to form candidate trips from m to n . Let $\text{dist}(m, n)$ be the network path length between m and n . For each sample, we compute the implied walking speed

$$v_i = \frac{\text{dist}(m, n)}{d_i^{\text{trav}}}, \quad (5)$$

discretize v_i into a speed category $v \in V$, and discretize t_i^{dep} into a time slot $t \in \mathcal{T}$ with slot length Δ . Aggregating sampled trips yields an initial demand scenario

$$S = \{D_{m,n,t,v} \mid m, n \in \mathcal{A}, t \in \mathcal{T}, v \in V\}, \quad (6)$$

where $D_{m,n,t,v}$ denotes the number of departures from m to n in time slot t with speed category v .

We then calibrate only the slot-wise absolute scale using counter observations:

$$r_t = \frac{\sum_{a \in \mathcal{A}^{\text{obs}}} L_{a,t}}{\sum_{a \in \mathcal{A}^{\text{obs}}} \sum_{m \in \mathcal{A}_a} \sum_{n \in \mathcal{A}} \sum_{v \in V} D_{m,n,t,v}}, \\ D'_{m,n,t,v} = \lceil r_t D_{m,n,t,v} \rceil, \quad (7)$$

where $\mathcal{A}_a \subseteq \mathcal{A}$ denotes areas whose induced flows can contribute to the count at area a under the adopted routing/assignment rule. This rescaling compensates for the mismatch between the absolute scale implied by $\{c_{m,n}\}$ and the observed counter volumes while preserving the relative OD/time pattern encoded by the prior. It matches the aggregate volume level per time slot; finer location-wise discrepancies are handled later by simulation and routing constraints.

From the refined scenario $S' = \{D'_{m,n,t,v}\}$, we construct a simulation-ready set of agent departure events. For each tuple (m, n, t, v) , we instantiate a total of $D'_{m,n,t,v}$ agents and assign each agent an origin–destination pair, a within-slot departure offset, and a walking-speed category consistent with v . This yields a network-feasible continuous-time initialization of the calibrated macroscopic scenario, which is then passed to the microscopic simulation layer described next.

3) *Trajectory Refinement via Multi-Agent Simulation*: The calibrated macroscopic scenario obtained in Sec. III-B2 provides a simulation-ready specification of area-to-area departures, but it is still coarse because it does not represent continuous pedestrian motion or local route choices within each area. We therefore refine the reconstructed area-level flows with a microscopic multi-agent simulation layer. The refined trajectories are then used to extract PoI-level decision events for training the destination-choice model.

To avoid confusion with the discrete time-slot index $t \in \mathcal{T}$, we use s to denote continuous simulation time. Let agent i have state $(\mathbf{x}_i(s), \mathbf{u}_i(s), g_i(s))$, where $\mathbf{x}_i(s) \in \mathbb{R}^2$ is position, $\mathbf{u}_i(s) \in \mathbb{R}^2$ is velocity, and $g_i(s)$ is the current navigation target. The kinematics are

$$\dot{\mathbf{x}}_i(s) = \mathbf{u}_i(s). \quad (8)$$

We adopt the Social Force model to incorporate goal-directed motion and interactions:

$$m_i \dot{\mathbf{u}}_i(s) = m_i \frac{u^{\text{max}} \mathbf{e}_i(s) - \mathbf{u}_i(s)}{\tau_i} + \sum_{j \neq i} \mathbf{f}_{ij}(s) + \sum_{o \in \mathcal{O}} \mathbf{f}_{io}(s), \quad (9)$$

where m_i is an inertial scaling, τ_i is a relaxation time, $\mathbf{e}_i(s)$ is the unit vector pointing toward $g_i(s)$, $\mathbf{f}_{ij}(s)$ is the interaction force from another agent j , and $\mathbf{f}_{io}(s)$ is the interaction force from obstacle $o \in \mathcal{O}$. Here, u^{max} denotes the simulator-assigned upper bound on the agent's admissible speed. This speed cap is determined by the current simulation context and may vary over time according to the agent's mobility mode. The dynamics are simulated with a fixed integration time step supported by the underlying simulator.

We implement this refinement by coupling SUMO [10] and JuPedSim [11]. SUMO enforces network feasibility (walkable topology, crossings, and signal control) and executes macroscopic routing, while JuPedSim provides microscopic pedestrian dynamics based on the Social Force model [11], [19]. This coupling yields continuous trajectories that remain consistent with the reconstructed area-level scenario while becoming locally plausible under congestion through interaction-aware motion.

To construct PoI-level decision data, we detect when a refined trajectory enters the vicinity of a PoI. Let $\mathcal{P}$ denote the set of PoIs used as candidate destinations and mobility endpoints. Each PoI $p \in \mathcal{P}$ has a spatial coordinate $\mathbf{r}_p \in \mathbb{R}^2$ and we regard agent i as visiting PoI p when

$$\|\mathbf{x}_i(s) - \mathbf{r}_p\| \leq \rho_p, \quad (10)$$

where ρ_p is the PoI-specific vicinity radius. At each such event, we record the current agent state, the local environment, and the next destination PoI.

C. Destination-Choice Model Design

The objective of the destination-choice model is to reproduce reconstructed movement trajectories. Because observational data are typically available for only one of the two conditions (with or without mobility introduction) under non-identical settings, it is difficult to directly measure how overall flows would change if a mobility service were introduced. Instead, we learn how visitors respond to environmental factors from observed data, and then treat mobility introduction as an environmental modification applied to the learned decision model.

Concretely, whenever an agent appears or arrives at a destination the decision model observes the current environment and the agent’s internal state, then selects the next destination. The process repeats until either the decision model chooses to exit or the agent’s stamina—which is decremented after every movement—reaches zero. The environmental inputs to the decision model include current time, PoI attractiveness values, and pairwise distances between PoIs. The agent-state inputs include the current location, previously visited PoIs, and cumulative walking distance

To implement this framework, we parametrize the decision model with neural networks and evaluate a set of candidate architectures. We consider two deterministic models: a multi-layer perceptron (MLP), which treats the input features as a fixed-length vector, and a graph neural network (GNN), which propagates information over the PoI graph through message passing. These variants are used to examine whether explicitly modeling PoI nodes and pairwise relations improves prediction accuracy. In contrast, we also consider a probabilistic variant, MLP + Mixture-of-Softmax (MoS), which learns a mixture distribution over candidate destinations and captures multi-modal choice behavior. During simulation, the next destination is sampled from this output distribution so that agents’ trajectories are stochastic and reflect decision uncertainty.

D. Online Phase: Mobility-Aware Simulation

In the online phase, agent behavior is generated by applying the trained destination-choice model within the simulator under a modified environment that reflects mobility introduction. Let f_Θ denote the trained decision model, and let $s_i(t)$ denote the current state of agent i at decision epoch t , and $\phi(t)$ denote the current environmental features. For deterministic models, $q_i^{next} = \arg \max_q f_\Theta(q | s_i(t), \phi(t))$; for probabilistic models, $q_i^{next} \sim f_\Theta(q | s_i(t), \phi(t))$.




Fig. 3: Example of Visualized Result.

To represent mobility introduction, we do not retrain the decision model on intervention-specific data. Instead, we modify the environmental features supplied to the trained model. Specifically, for mobility-enabled PoI pairs and affected PoI locations, we replace the effective travel speed with the mobility cruising speed and adjust the attractiveness of associated locations. The simulator-assigned upper bound on the agent’s admissible speed on link (p, q) is defined as

$$\tilde{u}_{p,q}^{max} = \begin{cases} u^{\max_{mob}}, & \text{mobility is available between } p \text{ and } q, \\ u^{\max_{walk}}, & \text{otherwise,} \end{cases} \quad (11)$$

and the corresponding pairwise distances between PoIs $\tilde{d}_{p,q}$ as

$$\tilde{d}_{p,q} = \frac{dist(p, q)}{\tilde{u}_{p,q}^{max}}. \quad (12)$$

Similarly, for locations affected by the introduction of mobility, attractiveness is updated.

These modified environmental variables are then fed into the trained decision model in place of the baseline features. As a result, the model reproduces how agents adapt their destination choices under mobility introduction without requiring a new intervention-specific training dataset. Operationally, agents are spawned according to the reconstructed departures, propagated on the network by SUMO/JuPedSim, and their targets or route segments are updated via TraCI whenever a new destination decision is triggered. This online procedure yields counterfactual agent trajectories under mobility introduction while preserving the macroscopic departure schedule estimated in the offline phase.

E. Visualization of Simulation Results

To facilitate intuitive understanding of mobility effects, we visualize simulated agent positions in three-dimensional space. During a SUMO run, TraCI enables sequential retrieval of each agent’s ID and location; by accumulating this time-series data, we can trace agents’ trajectories. We render these trajectories in Unity on top of Cesium’s 3D map tiles, placing agents according to the time-stamped positions obtained from TraCI. An example of the visualization output is shown in

Fig. 3; the rendered scene depicts Wakayama Castle Park, with agents rendered in yellow traversing the lower-right portion of the view. This visualization aids the interpretation of spatio-temporal changes in pedestrian flow by making local increases and decreases in movement visually apparent.

IV. EXPERIMENTAL EVALUATION

A. Objectives and Evaluation Metrics

We evaluate the proposed simulation platform by quantifying its ability to reproduce observed spatial and temporal patterns of pedestrian presence in the study area, using population distributions derived from field observations around Wakayama Castle Park. The evaluation pursues two principal aims: first, to assess the overall fidelity of simulated population distributions and their temporal dynamics; and second, to determine how specific modelling choices affect that fidelity. As primary quantitative metrics we use mean absolute error (MAE) computed across partitioned areas (and its day-aggregated sum) to measure spatial-count accuracy, and cosine similarity to assess agreement of temporal trajectories.

This study is organized around four research questions. First, to what extent the simulator reproduces observed spatial and temporal population distributions under mobility. Second, how the choice of destination decision model architecture, such as a multilayer perceptron or a graph neural network, influences reproduction quality and spatial error patterns, particularly in regions prone to spurious congestion. Third, whether probabilistic decision making improves aggregate prediction accuracy or uncertainty characterization compared with deterministic argmax selection. Fourth, which exit mechanism better reproduces empirical exit time distributions and their downstream effects on population dynamics.

B. Dataset and Preprocessing

Field observations were acquired within a 1 km² area centered on Wakayama Castle Park. Two sensing campaigns were carried out. The first campaign collected wide-area GPS location records continuously over one month (31 days) to capture typical, normal-state mobility patterns. These records were used as the wide-area location dataset for the GMM-based simulation procedure described in Sections III-B1 and III-B2, and to generate the baseline trajectory ensemble without mobility, which was used to construct the destination-choice training dataset. The second campaign was a separate seven-day field experiment in which high-resolution pedestrian spot counts were recorded with LiDAR sensors at preselected counting locations. These LiDAR measurements were used as spot-volume observations for calibrating the reconstructed flows, as described in Section III-B2. All spot measurements were aggregated into 10-minute intervals for downstream processing. All synthetic trajectories and training targets used in this work were generated using the method described in Section III-B1, III-B2, and III-B3. From the synthesized trajectory ensemble, we produce three derived data products that serve distinct purposes in the experiments.

TABLE I: PoI Attraction Values Without and With Mobility Introduction.

PoI ID	Area	Attraction Value	
		Without Mobility	With Mobility
00	A	0.0680	0.0545
01	A	0.0680	0.0545
02	D	0.157	0.151
03	G	0.138	0.153
04	F	0.140	0.152
05	C	0.00950	0.0990
06	B	0.0962	0.0978
07	E	0.0679	0.0628
08	H	0.165	0.176

Time- and location-specific arrival (spawn) rates and per-area, per-time-step population time series were generated for both the baseline (no mobility introduced) and the mobility-introduced conditions — the mobility-introduced versions were obtained by calibrating reconstructed flows to the seven days of LiDAR observations collected during the intervention week. The destination-choice training dataset (decision-epoch records paired with the selected next destination, including state and spatial features such as current PoI, time of day, cumulative travel distance, and distances to candidate PoIs) was created only from the baseline (without mobility) reconstructed trajectories and used to train the decision models under normal-state behavior assumptions.

C. Simulation Settings

We performed simulations covering a single day (86,400 seconds, 24 hours) for Wakayama Castle Park in Japan. The park geometry and road network were reconstructed from OpenStreetMap (OSM) data.

1) *Area Partitioning and Attraction Settings*: The park area was partitioned into areas used to define each pedestrian's origin and destination areas within the simulation; population distributions are computed with respect to these areas. We defined specific landmarks (e.g., the zoo, the castle keep, parking lots) as PoIs that are expected to attract frequent visits. Each area contains at least one PoI. Visualizations of the area partitioning and PoI locations are shown in Fig. 5 and Fig. 6. Each PoI is assigned an attraction score representing the propensity of visitors to want to visit that location. Attraction scores were derived from the per-area visit proportions in the dataset described in Section IV-B and normalized across all PoIs. If multiple PoIs are contained within a single area, the area's visit proportion is split equally among those PoIs. Because the attraction scores are normalized over all PoIs, the score of a given location may decrease in relative terms with mobility introduction as the scores of other locations increase. The mapping from PoIs to zones and the attraction values without/with mobility introduction used in the experiments are given in Table I.

2) *Mobility Introduction Settings*: In accordance with the field experiment, mobility was introduced on the routes indicated in green in Fig. 6. For the experiments, we set the mobility vehicle cruising speed to 20 km/h and the pedestrian



(a) C⁺walk_T
three-wheeled standing-type battery electric vehicle (BEV).¹



(b) C⁺walk_S
pedestrian-mobility-assistance model.²



(c) Kintone Model One S
electric kick scooter.³



(d) Yamaha AR-07 Green Slow Mobility
low-speed electric vehicle.⁴

Fig. 4: Mobilities we installed.




Fig. 5: Area Partitioning for our Simulation.




Fig. 6: PoI Locations and Mobility Deployment.

walking speed to 5.0 km/h. In addition, attraction scores were updated at deployment locations such as location_00, location_01, location_03, and location_04. The micromobilities we

TABLE II: Population Reproduction Accuracy With Mobility Introduction.

Model	MAE	Day-aggregated MAE
MLP	2.34	334
MLP + MoS	2.52	360
GNN	2.87	410
MLP with Exit Class	1.76	252
MLP + MoS with Exit Class	1.75	250

TABLE III: Cosine similarities for two evaluation criteria.

Model	Population	Change in Population
MLP	0.769	0.664
MLP+MoS	0.758	0.694
GNN	0.663	0.578
MLP with Exit Class	0.624	0.582
MLP + MoS with Exit Class	0.653	0.619

installed are shown in Fig 4.

D. Reproducing Human Flows With Mobility Introduction

1) *Overall Fidelity of the Proposed Platform:* We first evaluate the overall fidelity of the proposed simulation platform in reproducing observed pedestrian flows. Figure 7 presents the simulated temporal population distributions per area together with the ground-truth distribution. Table II summarizes the MAE and day-aggregated MAE for each model configuration, while Table III reports two cosine-similarity metrics: the similarity between simulated and observed population time-series under the mobility-enabled scenario, and the similarity of the mobility-induced change patterns computed as the difference between the mobility-enabled and baseline time-series.

Regarding spatial population reproduction, the day-aggregated MAE ranges from approximately 250 to 400 depending on the model configuration. Across all models, Regions F and G tend to be overestimated, suggesting that the simulator may slightly overpredict the increase in visitor concentration caused by the mobility introduction. Nevertheless, the cosine similarity results indicate that the simulator

¹Toyota Motor Corporation, “Toyota Launches the C+walk T in Japan, a New Form of Walking-Area Mobility,” <https://global.toyota/en/newsroom/toyota/35872225.html>.

²Toyota Motor Corporation, “Toyota Launches the C+walk S in Japan, a New Form of Walking-Assistance Mobility,” <https://global.toyota/en/newsroom/toyota/38947604.html>.

³Kintone official product page, “Kintone Model One S,” <https://kintone.mobi/products/modelone-s?variant=42886564970678>.

⁴Yamaha Motor Co., Ltd., “Low-Speed Electric Mobility for First- and Last-Mile Transportation,” <https://global.yamaha-motor.com/news/2023/0825/newsletter.html>.




Fig. 7: Estimated and Ground Truth Population Distributions With Mobility Introduction.

captures the overall temporal patterns of population dynamics reasonably well. For both the population trajectories and the mobility-induced changes, cosine similarities are ranging from 0.578 to 0.769, with the MLP reaching 0.769, indicating that the platform can reproduce the general temporal trends of pedestrian flows in the study area.

2) *Effect of Network Architecture*: Next, we investigate how the architecture of the destination-decision model affects the reproduction of human flows by comparing MLP and GNN models. As shown in Figure 7, only the GNN model produces a spurious congestion pattern in Area E that does not appear in the ground-truth data. This behavior is also reflected in the larger MAE values reported in Table II.

These results suggest that explicitly modeling spatial relationships between PoIs through a graph structure does not necessarily improve aggregate flow reproduction in this setting. Because the study area covers only approximately 1 km^2 , the spatial distances between areas are relatively small and walking remains a feasible option across most locations. Consequently, the benefits of modeling detailed spatial relationships through a GNN appear limited in this compact environment.

3) *Probabilistic vs. Deterministic Decision Modeling*: We further examine the effect of probabilistic decision modeling. Specifically, we compare a deterministic MLP model, which selects the next destination using an argmax rule, with an MLP augmented with a Mixture-of-Softmax (MoS) output layer that learns a distribution over candidate destinations and samples the next destination according to that distribution.

The results shown in Table II and Table III indicate that

there is no substantial difference in accuracy between the two approaches. Both the MAE of the reproduced population distribution and the cosine similarity of the temporal population trajectories are comparable between the deterministic and probabilistic models. This suggests that, under the present experimental conditions, modeling decision uncertainty through a probabilistic output distribution does not significantly improve the aggregate reproduction accuracy of pedestrian flows.

4) *Exit Mechanism and Its Impact*: Next, we compare two approaches for modeling visitors' exit behavior from the park. In the first approach, exiting the park is treated as an explicit class within the destination-choice model, allowing the exit timing to be learned directly as part of the classification task. In the second approach, exit behavior is modeled using a stamina variable that decreases as the agent moves; agents leave the simulation when the stamina value reaches zero, and the stamina parameters are tuned based on the training data.

In Figure 7, Table II, and Table III, the models labeled with Exit Class correspond to the former approach, while the remaining models employ the stamina-based mechanism. Table II shows that incorporating the exit event as an explicit class leads to lower MAE values, indicating improved accuracy in reproducing the spatial population distribution. In contrast, the cosine similarity results in Table III suggest that the stamina-based mechanism achieves slightly better agreement with the observed temporal population trajectories. These results indicate that each approach has distinct advantages, and the appropriate choice may depend on the evaluation objective of the simulation.




Fig. 8: Grouped SHAP beeswarm plot averaged over classes.

5) *Analysis of the Decision-making Model:* Finally, Fig. 8 shows the grouped SHAP summary (beeswarm) plot for the MLP-based decision model [28]. Here, features are organized into groups: the *current PoI* group represents a set of variables indicating the agent’s current PoI, and the *visited PoI* group represents variables indicating whether each PoI has been previously visited. Each point represents one sample, and the horizontal position indicates the SHAP value (i.e., the contribution to the model output), where positive and negative values respectively increase or decrease the predicted probability. Colors indicate low-to-high feature values within each group. Across samples, the *current PoI* shows the widest spread of SHAP values, indicating that it is the most influential group and can substantially increase or decrease the model output depending on the context. The *visited PoI* group also exhibits a broad distribution on both sides of zero, suggesting heterogeneous and potentially non-monotonic effects across samples (i.e., higher group values do not consistently correspond to either positive or negative contributions). In contrast, *total travel distance* and *time of the day* show smaller magnitudes concentrated near zero, implying more moderate, context-dependent contributions. The *distance between PoI* group remains tightly centered around zero, indicating a negligible effect on the model output under the present data and grouping scheme. Although inter-PoI distance appears to be an important decision-making feature, our observations for Wakayama Castle suggest otherwise: because the distances between areas are, in practice, short enough that walking does not impose a burden, inter-PoI distance is unlikely to constitute a salient factor in visitors’ decision-making.

V. CONCLUSION

We presented a platform that predicts and visualizes the effects of introducing mobility services to improve visitor circulation, based on a human-flow digital twin. The digital twin integrates a multi-agent simulator capable of representing destination choice as a function of agents’ current state and surrounding environmental information. For agent-level decision making, we implemented an MLP and a GNN model tuned by observed decision instances.

We evaluated the proposed framework using human-flow data collected in Wakayama Castle Park in Japan. The MLP model reproduced the temporal evolution of area-level population distributions with mobility introduction with cosine similarities exceeding 0.7, indicating that the platform captures the overall trend of population redistribution induced by mobility introduction. Furthermore, the MLP reproduced the change in population distribution caused by mobility introduction with a cosine similarity of 0.664, showing its ability to reflect mobility-induced behavioral shifts.

For future work, we will first focus on improving the fidelity of the decision-making model, particularly by explicitly incorporating waiting actions into the model structure. Furthermore, we will extend the platform’s visualization capabilities to support real-time rendering of simulation outcomes.

REFERENCES

- [1] H. Badia and E. Jenelius, “Shared e-scooter micromobility: review of use patterns, perceptions and environmental impacts,” *Transport Reviews*, vol. 43, no. 5, pp. 811–837, 2023.
- [2] Y. Yang, L. Jiang, and Z. Zhang, “Tourists on shared bikes: Can bike-sharing boost attraction demand?” *Tourism Management*, vol. 86, p. 104328, 2021.
- [3] D. Kazyeva, P. Stutz, G. Wallentin, and M. Loidl, “Large-scale agent-based simulation model of pedestrian traffic flows,” *Computers, Environment and Urban Systems*, vol. 105, p. 102021, 2023.
- [4] S. Hörnl, “Integrating discrete choice models with matsum scoring,” *Procedia Computer Science*, vol. 184, pp. 704–711, 2021.
- [5] M. Dozza, A. Violin, and A. Rasch, “A data-driven framework for the safe integration of micro-mobility into the transport system: Comparing bicycles and e-scooters in field trials,” *Journal of safety research*, vol. 81, pp. 67–77, 2022.
- [6] M. Mishra, S. M. Kala, T. Amano, and H. Yamaguchi, “Balancing privacy and planning: using counterfactuals to predict and optimize tourism in wakayama city,” in *Proceedings of the 2nd ACM SIGSPATIAL international workshop on geo-privacy and data utility for smart societies*, 2024, pp. 25–30.
- [7] R. Naeem, S. Manas, T. Amano, and H. Yamaguchi, “A questionnaire-only counterfactual machine learning approach to assess the spatial impact of green mobility vehicles in urban parks: A wakayama castle case study,” in *Companion Proceedings of the 27th International Conference on Distributed Computing and Networking*, ser. ICDCN Companion ’26. Association for Computing Machinery, 2026, p. 120–125.
- [8] D. Nag, F. Brandel-Tanis, Z. A. Pramestri, K. Pitera, and Y. K. Frøyen, “Exploring digital twins for transport planning: a review,” *European Transport Research Review*, vol. 17, no. 1, p. 15, 2025.
- [9] M. Uegaki, T. Amano, and H. Yamaguchi, “Simulating urban pedestrian flows by fusing wide-area location data and spot pedestrian counts,” in *Proceedings of the 21st EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2024.
- [10] C. R. Daniel Krajzewicz, Georg Hertkorn and P. Wagner, “Sumo (simulation of urban mobility): an open-source traffic simulation,” in *Proceedings of 4th Middle East Symposium on Simulation and Modelling*, 2002.
- [11] J. Z. Arnel Ulrich Kemloh Wagoum, Mohcine Chraïbi and G. Lamme, “Jupedsim: an open framework for simulating and analyzing the dynamics of pedestrians,” in *Proceedings of the 3rd Conference of Transportation Research Group of India*, 2015.
- [12] A. Wegener, M. Piorkowski, M. Raya, H. Hellbrück, S. Fischer, and J.-P. Hubaux, “Traci: An interface for coupling road traffic and network simulators,” in *Proceedings of the 11th Communications and Networking Simulation Symposium*, April 2008.
- [13] T. Robin, G. Antonini, M. Bierlaire, and J. Cruz, “Specification, estimation and validation of a pedestrian walking behavior model,” *Transportation Research Part B: Methodological*, vol. 43, no. 1, pp. 36–56, 2009.

- [14] J. Gao, P. Peng, F. Lu, C. Claramunt, and Y. Xu, "Towards travel recommendation interpretability: Disentangling tourist decision-making process via knowledge graph," *Information Processing & Management*, vol. 60, no. 4, p. 103369, 2023.
- [15] F. Dai, D. Wang, and K. Kirillova, "Travel inspiration in tourist decision making," *Tourism Management*, vol. 90, p. 104484, 2022.
- [16] S. McCabe, C. Li, and Z. Chen, "Time for a radical reappraisal of tourist decision making? toward a new conceptual model," *Journal of Travel Research*, vol. 55, 07 2015.
- [17] C. King and N. W. Bode, "A virtual experiment on pedestrian destination choice: the role of schedules, the environment and behavioural categories," *Royal Society open science*, vol. 9, no. 7, p. 211982, 2022.
- [18] C.-J. Jin, Y. Luo, C. Wu, Y. Song, and D. Li, "Exploring the pedestrian route choice behaviors by machine learning models," *ISPRS International Journal of Geo-Information*, vol. 13, no. 5, p. 146, 2024.
- [19] D. Helbing, "A mathematical model for the behavior of pedestrians," *Behavioral science*, vol. 36, no. 4, pp. 298–310, 1991.
- [20] A. Treuille, S. Cooper, and Z. Popović, "Continuum crowds," *ACM transactions on graphics (TOG)*, vol. 25, no. 3, pp. 1160–1168, 2006.
- [21] S. Zhang, J. Zhang, M. Chraïbi, and W. Song, "A speed-based model for crowd simulation considering walking preferences," *Communications in nonlinear science and numerical simulation*, vol. 95, p. 105624, 2021.
- [22] H.-T. Dang, B. Gaudou, and N. Verstaëvel, "Hypedsim: A multi-level crowd-simulation framework—methodology, calibration, and validation," *Sensors*, vol. 24, no. 5, p. 1639, 2024.
- [23] R. Huang, X. Zhao, Y. Yuan, Q. Yu, C. Liu, and W. Daamen, "Modeling pedestrian tactical and operational decisions under risk and uncertainty: A two-layer model framework," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 5, pp. 5259–5281, 2023.
- [24] A. Horni, K. Nagel, and K. W. Axhausen, "High-resolution destination choice in agent-based demand models," *Arbeitsberichte Verkehrs-und Raumplanung*, vol. 682, 2011.
- [25] T. Wörle, L. Briem, M. Heilig, M. Kagerbauer, and P. Vortisch, "Modeling intermodal travel behavior in an agent-based travel demand model," *Procedia Computer Science*, vol. 184, pp. 202–209, 2021.
- [26] H. Yang, E. Wong, H. Davis III, and J. Y. Chow, "A co-simulation system that integrates matsim with a set of external fleet simulations," *Simulation Modelling Practice and Theory*, vol. 134, p. 102957, 2024.
- [27] J. Li, E. Rombaut, and L. Vanhaverbeke, "Agent-based digital traffic model generation for regions facing data scarcity using aggregated cellphone data: a case study for brussels," *International Journal of Digital Earth*, vol. 17, no. 1, p. 2407046, 2024.
- [28] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in neural information processing systems*, vol. 30, 2017.