Integrating Preference-Aware Modeling of Human Spatial Behavior in Cyber-Physical-Human Systems

Maral Doctorarastoo[1] mdoctora@andrew.cmu.edu

Katherine A. Flanigan[1]* kflaniga@andrew.cmu.edu

Mario Bergés[1] mberges@andrew.cmu.edu

Christopher McComb[2] ccm@andrew.cmu.edu

[1]**Department of Civil and Environmental Engineering**
[2] **Department of Mechanical Engineering**
Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

This study introduces a new approach for modeling preference-aware human spatial behavior using Graph Neural Networks (GNN) and Reinforcement Learning (RL). Current models often overlook the causality and impact of factors influencing preferences. Our approach utilizes GNN for its advanced handling of graph-structured spatial data, capturing physical, social, and environmental features and how these are perceived by humans. Integrated with RL, the model dynamically adapts to changes in the surrounding environment, improving adaptability and generalizability of simulations. As a proof of concept, we illustrate the approach in an educational conference room setting to compare student behavior simulation with and without preference inclusion. The results indicate that preference incorporation leads to significantly more realistic simulations, highlighting its potential to improve the design and control of cyber-physical-human systems.

1 Introduction

In recent years, advancements in cyber-physical-human systems (CPHS) have focused on physical simulations, feedback algorithms, and sensor integration [13]. A key challenge remains in characterizing human models and understanding how they adapt during interactions [1]. This paper focuses on cognitive-level human behavioral interactions, treating humans as decision makers who interact with cyber-physical systems through spatio-temporal behaviors. The primary challenge lies in developing computational models of human spatial behavior, which are essential for system design and control in applications such as human-robot collaboration, self-driving vehicles, and smart environments.

Various simulation approaches, including system dynamics, flow-based, and multi-agent simulation (MAS) [15], have been used to model human behavior in the built environment at different levels of abstraction. Among these, MAS is notable for offering simulations that capture complex agent interactions, enabling autonomous decision making and modeling emergent behaviors. However, defining the rules that govern agents' interactions with each other and their changing environments is challenging due to the range of factors that influence behavior and preferences. These include physiological (e.g., age, gender, sensory perceptions), psychological (e.g., motivations, needs, preferences), spatial and environmental (e.g., temperature, proximity to others, seat arrangements), time-related (e.g., time of day), contextual (e.g., social constraints, routines), social ties, and random factors.

*Corresponding author.

NeurIPS 2024 Workshop on Behavioral Machine Learning.

Researchers have long sought to develop the rules that govern spatial behavior by considering various influential factors. These models generally fall into two categories: *knowledge-driven*, which are based on survey data, observational studies, and expert insights [3], and *data-driven* [18], which rely on extensive datasets enabled by advances in the Internet of Things and sensor technology [11, 10]. The training data for these models includes both known and unknown variables that can influence decisions. However, these models often struggle to generalize to environments outside their training data. Hybrid models, such as those using Reinforcement Learning (RL), combine the strengths of different modeling approaches. RL excels in dynamic learning, as it can handle complexity and uncertainty in evolving environments. RL also mirrors human decision-making processes by being goal-driven, sequential, and focused on maximizing internal utility functions of agents.

A major challenge in modeling human behavior is the difficulty of defining the internal utility function that represents individual goals and preferences. Developing this function based on data is complex, not only because of the individual-specific factors influencing human decisions but also because the human agent must perceive relational configurations. This involves understanding the spatial arrangement and relative positioning of elements and other individuals within the environment. By accurately modeling these spatio-temporal perceptions, control actions can be implemented to encourage desired behaviors, thereby aligning with design objectives.

To address the complex needs of accounting for factors that influence human decisions *and* the perception of relational configurations within their environment, we use Graph Neural Networks (GNN) in our study. GNNs are uniquely capable of capturing, processing, and aggregating relational data [14, 16], making them particularly effective in representing physical, social, and environmental features of an environment and their interconnections. We use GNN as a key component to map complex human perceptions and preferences to reward functions. The proposed GNN-based preference model enhances the realism of the reward function in the RL algorithm, compared to simply allocating rewards for performing activities. This approach results in behavioral simulations that are more responsive to the diverse physical, social, and environmental attributes of the built environment.

We first discuss the development of the GNN-based preference model, followed by how RL is used to model human spatial behavior in dynamic environments. We then describe the integration of GNN and RL to create a preference-aware model. Finally, we demonstrate the model's practical applications through two conference room scenarios, highlighting improvements in simulation realism and behavioral prediction compared to non-preference-based approaches.

2 Methodology

2.1 GNN preference model

Data collection. Accurate, *in-situ* data collection is crucial for building a reliable GNN-based preference model. In public settings, privacy-preserving technologies like the Azure Kinect are used to generate data processed by activity recognition algorithms [11, 10]. Such datasets can be further enhanced with tracklets that capture precise activity locations, while a social map monitors individual movements over time. Sensors track ambient conditions, and additional layers, such as time stamps and the physical layout (via digital blueprints or LiDAR), contribute a comprehensive dataset. This rich data foundation serves as a key input for the GNN preference model.

Graph representation of the environment. To apply GNN, the collected dataset must be structured as a graph. The environment is divided into a grid, with each cell mapped to a node in the graph $G = (V, E)$, where V represents nodes, and E represents edges connecting adjacent cells. The adjacency matrix A is defined as $a_{ij} = 1$ if cells i and j are adjacent, and $a_{ij} = 0$ otherwise. Each node $v_i \in V$ is associated with a feature vector $x_i = [Er_i, En_i, S_i]$, where Er_i represents ergonomic features, En_i reflects environmental attributes, and S_i indicates social occupancy. These features are context-specific; for instance, in educational settings, they may include elements like furniture, temperature, and light. Our approach utilizes node classification within the GNN model to estimate the probability of selecting each grid cell (node) for specific activities. The target vector then assigns labels indicating activity selection for each node.

Implementation of the preference model. The preference model operates in two phases: *Training* and *Prediction*. During *Training*, the model calculates the probability of each grid cell being selected for activities based on spatial, environmental, and social features from the *Input Graph*. This selection

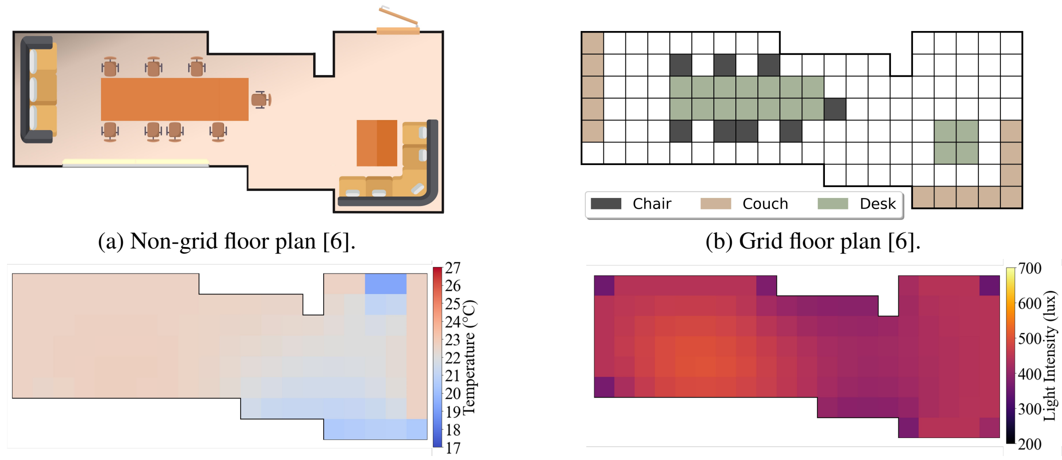

(a) Non-grid floor plan [6].

(b) Grid floor plan [6].

(c) Temperature heatmap.

(d) Light intensity heatmap.

Figure 1: Overview of the environment layout and heatmaps.

process is influenced by the dynamics of human physiological and psychological traits, as well as preferences related to the relational configuration of physical space, all captured through sensor data. The *Output Graph* reflects cell selection, guided by these human preferences. To address data imbalance caused by selecting only one node from the entire graph, we apply a weighted Binary Cross-Entropy (BCE) loss function. This function assigns greater weight to the underrepresented class (positive samples), enhancing their impact on the loss calculations and reducing the bias towards the majority class. This adjustment is crucial in applications like preference modeling where the minority class is of greater interest. After training, the model enters the *Prediction* phase, where it calculates the probability of selecting each grid cell for each possible activities.

2.2 Preference-aware RL spatial behavior model

Specification of RL behavior model. In our model, humans are represented as RL agents whose spatial decision making is guided by policy functions. The environment includes the cyber, physical, and social worlds in which agents operate. Agents observe the state of the environment, which may include only partial information from the surrounding. Based on these observations they perform specific activities through fine-grained actions such as moving, sitting, standing, and initiating interactions. The agent's goal is to successfully accomplish these activities, for which it receives rewards. However, this basic reward function overlooks spatial preferences, leading to homogeneous behaviors that fail to accurately reflect human decision making. The next section addresses this issue.

GNN integration with RL. The integration process begins with the RL agent observing the current state of the environment. Instead of using a simplistic reward function, the agent uses the GNN's predicted likelihood of various spaces being chosen for different activities given the current state of the environment. This allows the agent to assess the desirability of different spatial locations for performing activities. The rewards the agent receives are thus based on its complex perceptions and spatial preferences as predicted by the GNN. At each time step, the agent's policy function evaluates possible actions based on the GNN-predicted rewards, which indicate the suitability of locations. This results in agents choosing locations that align with their specific preferences. For instance, if the GNN predicts a high likelihood that a location is ideal for studying, the agent receives a higher reward for choosing it. These rewards can change as the physical and social characteristics of the surroundings evolve, potentially prompting the agent to relocate if the suitability of a location changes significantly, leading to more human-like spatial behaviors.

3 Illustrative example

A practical application of human spatial behavior models is in cyber-physical-social infrastructure systems, where infrastructure is designed and controlled to meet economic and human-centered objectives [4, 5]. To illustrate, we develop a human preference model in an educational facility.

Case study setting. The conference room, shown in Figure 1a, is a 4m$\times$10m space that has been divided into a grid (Figure 1b), converting the room into a discrete spatial model for graph representation. Each 50cm$\times$50cm grid cell represents physical elements like walls, entrances, windows, furniture, and unoccupied spaces. The main activities are studying, eating, and socializing.

Graph structure. A *mesh* is extracted from the grid, excluding unoccupiable cells and walls, focusing on potential areas for human presence. A graph overlays the mesh, where nodes represent mesh cells and edges connect adjacent nodes, including diagonals. The feature vector of each node includes environmental features (e.g., light, temperature), the objects presence (e.g., tables, chairs), proximity to architectural elements (e.g., windows, walls), and occupancy.

Creating the synthetic dataset. For this proof-of-concept study, we generated 10,000 synthetic data points per activity to analyze spatial preferences. Node feature vectors were constructed based on room characteristics, while social and environmental maps were created using probabilistic distributions to better mirror real-world conditions and enrich the dataset. The social map, representing individual locations, is created by randomly placing individuals to simulate the presence of others, affecting preferences for social interactions or quiet zones suitable for focused tasks. Light intensity follows a normal distribution [2] with an average of 600 lux near windows, 300 lux in corners, and 500 lux under artificial lights, with corresponding standard deviations of 200, 150, and 75 lux. Similarly, temperature is modeled using a normal distribution [12], averaging 26°C near heating/windows, 22°C in corners, and 19°C near the entrance, with respective standard deviations of 3°C, 1.5°C, and 2°C. Nearest-neighbor interpolation is used for light [17], and cubic interpolation for temperature [8].

Human activity preferences are modeled through a probabilistic selection mechanism that evaluates cells for suitability as activity zones. Suitability is calculated using a weighted sum of node features, with weights representing the importance of each feature for a given activity, influencing levels of preference, indifference, or avoidance. For example, someone studying might prioritize light intensity and proximity to windows, whereas for socializing, they may prioritize furniture layout. Each data instance includes selecting an initial node for the activity's starting point and a target node chosen based on feature-weighted preferences, reflecting preferences within the current spatial, social, and environmental context. In this case study, three human agents are placed in the conference room. *Agent 1*, while studying, prefers areas with couches, avoids high temperatures, and is indifferent to other factors. *Agent 2* prefers warmer areas, avoids crowded spaces, and is indifferent to other conditions. *Agent 3* seeks well-lit areas near windows and is indifferent to other features.

Preference model training. The GNN is trained iteratively using the Adam optimizer (learning rate 0.0005) to minimize prediction errors, with early stopping after 50 epochs to prevent overfitting. The model uses a batch size of 128, ReLU and Sigmoid activation functions for hidden and output layer respectively, and a weighted BCE loss function. The learning rate is scheduled with a step size of 50 and a decay factor, γ, of 0.99, with a maximum of 1000 epochs. The downward trend in training and test loss curves (Figure 2) and the convergence of these curves indicates a balanced adaptation to the training data while maintaining generalization to new, unseen data.

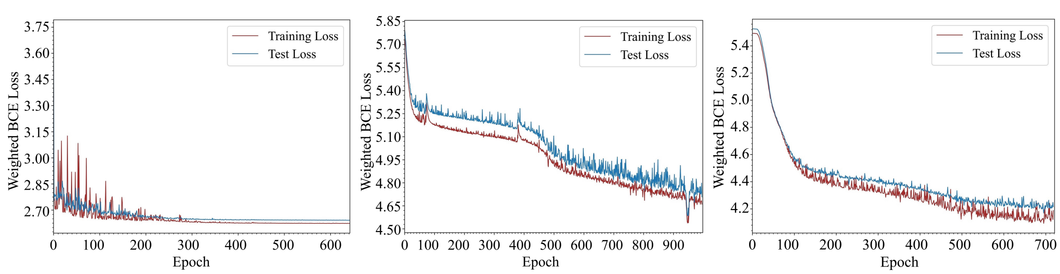

(a) (b) (c)

Figure 2: GNN training and test loss over epochs for (a) *Agent 1*, (b) *Agent 2*, and (c) *Agent 3*.

4 Results and conclusion

We simulated movement trajectories and usage patterns for three agents using agent-based simulation. Figures 1c and 1d display temperature and light maps, while Figures 3a, 3b, and 3c compare agent behavior with and without preference modeling across 200 trajectories. Without preference modeling, agents typically choose the nearest seating to the entrance, reflecting a basic utilitarian approach where minimizing travel distance is the primary goal. The GNN-based preference model significantly alters these behaviors, tailoring them to each agent's preferences and leading to more complex trajectories. When agents operate under preference-aware decision-making, they often traverse longer paths to reach their desired locations, suggesting that such preference-driven actions promote goal-oriented movement, albeit with increased travel distances. These results underscore the relevance of preference-based modeling for applications like urban planning and smart environments, where user-specific goals can inform design optimizations. *Agent 1*, who prefers couches and cooler areas, predominantly selects the couch in the cooler part of the room, prioritizing thermal comfort over proximity to the entrance. *Agent 2*, who seeks warmth and solitude, moves toward a couch on the far-left side, balancing warmth with social distance. The presence of other agents, particularly *Agent 3*, influences this choice, highlighting the multi-agent interaction in the simulation. As seen in Figure 3c, *Agent 3*—who prefers well-lit seating near windows—consistently selects spots close to the window. This choice encourages *Agent 2* to prioritize the couch over nearby chairs, despite similar temperatures. The results demonstrate that integrating preferences leads to behaviors that more closely mirror human decision making, resulting in more diverse and realistic movements within the space compared to the simple behavior observed without preference modeling.

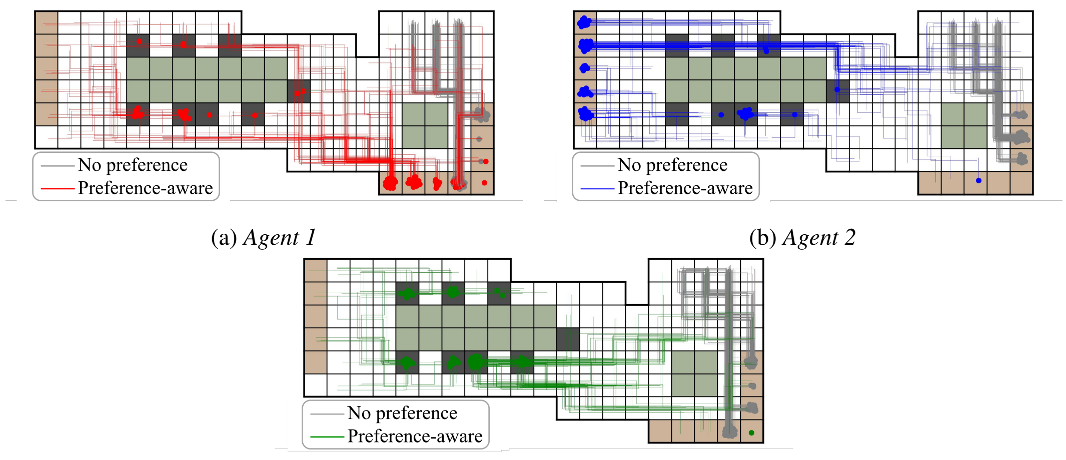

(a) *Agent 1* (b) *Agent 2*

(c) *Agent 3*

Figure 3: Comparison of agent trajectories under preference-aware and no preference conditions.

This study represents an initial effort to address a grand challenge of CPHS: modeling and predicting human behavior. It lays the groundwork for innovative applications that integrate sensing technology and artificial intelligence with human-centric design principles, promising a future where physical systems adapt seamlessly to human needs. Future work should focus on enabling near real-time model retraining to adapt to dynamic changes in preferences and goals. Additionally, assessing the representativeness of training data concerning human physiological and psychological characteristics, alongside benchmark comparisons, remains essential. Further model enhancements could include integrating internal factors such as emotions [9] and social ties [7], refining the reward discount factor to balance immediate and future objectives, and validating these models with real-world data to ensure their practical applicability.

Acknowledgments and Disclosure of Funding

This material is based upon work supported by the National Science Foundation under Grant Number 2425121 and the Carnegie Mellon University Rankin Fellowship. Mario Bergés holds concurrent appointments as a Professor of Civil and Environmental Engineering at Carnegie Mellon University and as an Amazon Scholar. This publication describes work performed at Carnegie Mellon University and is not associated with Amazon.

References

[1] A. M. Annaswamy, K. H. Johansson, and G. J. Pappas. *Control for Societal-scale Challenges: Road Map 2030*. IEEE Control Systems Society, 2023.

[2] Y. Cohen, R. Iribarren, H. Ben-Eli, A. Massarwa, N. Shama-Bakri, and O. Chassid. Light intensity in nursery schools: a possible factor in refractive development. *Asia-Pacific Journal of Ophthalmology*, 11(1):66–71, 2022.

[3] M. J. Dawes and M. J. Ostwald. Testing the 'wright space': using isovists to analyse prospect-refuge characteristics in usonian architecture. *The Journal of Architecture*, 19(5):645–666, 2014.

[4] M. Doctorarastoo, K. A. Flanigan, M. Bergés, and C. McComb. Exploring the potentials and challenges of cyber-physical-social infrastructure systems for achieving human-centered objectives. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '23*, page 385–389, Istanbul, Turkey, 2023. doi: 10.1145/3600100.3626340.

[5] M. Doctorarastoo, K. A. Flanigan, M. Bergés, and C. McComb. Modeling human behavior in cyber-physical-social infrastructure systems. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '23*, page 370–376, Istanbul, Turkey, 2023. doi: 10.1145/3600100.3626338.

[6] M. Doctorarastoo, K. A. Flanigan, M. Bergés, and C. McComb. GNN-based predictive modeling of human preferences in the built environment. In *ASCE International Conference on Computing in Civil Engineering 2024 (i3ce 2024)*, Pittsburgh, PA, USA., 2024.

[7] M. Doctorarastoo, C. Lin, K. Flanigan, and M. Bergés. A Bayesian network approach for predicting social interactions in shared spatial environments. In *Proceedings of the 11th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '24*, Hangzhou, China, 2024. doi: 10.1145/3671127.3699530.

[8] H. Li, J. Vink, and F. Alpak. An efficient multiscale method for the simulation of in-situ conversion processes. *SPE Journal*, 20(03):579–593, 2014. doi: 10.2118/172498-PA.

[9] C. Lin, M. Doctorarastoo, and K. Flanigan. Your actions talk: automated sociometric analysis using kinesics in human activities. In *Proceedings of the 11th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '24*, Hangzhou, China, 2024. doi: 10.1145/3671127.3699529.

[10] C. Lin, J. Martins, and K. A. Flanigan. Read the room: inferring social context through dyadic interaction recognition in cyber-physical-social infrastructure systems. In *ASCE International Conference on Computing in Civil Engineering 2024 (i3ce 2024)*, Pittsburgh, PA, USA., 2024.

[11] J. Martins, K. Flanigan, and C. McComb. Skeleton-based human action recognition in a thermal comfort context. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 377–384, Istanbul Turkey, 2023. doi: 10.1145/3600100.3626339.

[12] S. Qiu, T. Asyhari, W. Guo, S. Wang, B. Li, C. Zhao, and M. Leeson. Molecular channel fading due to diffusivity fluctuations. *IEEE Communications Letters*, 21(3), 2017. doi: 10.1109/LCOMM.2016.2633988.

[13] T. Samad. *Human-in-the-Loop Control and Cyber–Physical–Human Systems: Applications and Categorization*, pages 1–23. IEEE Press, Piscataway, NJ, USA, 2023.

[14] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

[15] D. Schaumann and M. Kapadia. Modeling social and spatial behavior in built environments: current methods and future directions. *Social-Behavioral Modeling for Complex Systems*, pages 673–695, 2019.

[16] M. Taghizadeh, K. Khayambashi, M. A. Hasnat, and N. Alemazkoor. Multi-fidelity graph neural networks for efficient power flow analysis under high-dimensional demand and renewable generation uncertainty. *Electric Power Systems Research*, 237:111014, 2024.

[17] Z. Zhao, J. Wang, X. Zhao, C. Peng, Q. Guo, and B. Wu. Navilight: indoor localization and navigation under arbitrary lights. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.

[18] R. Zhu, B. Becerik-Gerber, J. Lin, and N. Li. Behavioral, data-driven, agent-based evacuation simulation for building safety design using machine learning and discrete choice models. *Advanced Engineering Informatics*, 55:101827, 2023.