Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 58-30 (2024) 115–120

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

Maral Doctorarastoo* Katherine A. Flanigan†* Mario Bergés*

* *Department of Civil and Environmental Engineering, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, (e-mails: mdoctora@andrew.cmu.edu, kflaniga@andrew.cmu.edu,mberges@andrew.cmu.edu)*

Abstract: This study introduces a new approach for modeling preference-aware human spatial behavior in cyber-physical-human systems 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 spatial data, capturing physical, social, and environmental features and their human perception. Integrated with RL, the model dynamically adapts to changes in the surrounding environment. We illustrate the approach in an educational conference room setting, comparing student behavior simulations with and without preferences. The results indicate that preference incorporation leads to significantly more realistic simulations.

Keywords: Agent-based modeling, cyber-physical-human systems, human behavior modeling, human-infrastructure interaction, preference modeling.

1. INTRODUCTION

Cyber-physical-human systems (CPHS) integrate control systems with human interaction, enabling adaptive interaction across various levels: physiological, cognitive-behavioral, and population-wide levels. Notable progress has been made in the study of CPHS, particularly in technological areas such as physical world simulations, feedback algorithms, and sensor integration (Samad (2023)). However, despite notable differences in modeling, analysis, control objectives, and design tools across the various levels of CPHS, common central challenges remain to be addressed. Due to the complexity and often limited access to measurements, human models are inherently highly uncertain. Understanding how these human models adapt during interaction remains a critical research goal (Annaswamy et al. (2023)).

In this paper, we focus on behavioral interactions at the cognitive-behavioral level, where humans are viewed as decision makers who interact with cyber-physical systems (CPS) in a spatio-temporal fashion via behaviors. Here, "spatial" indicates that there are observable actions and reactions that individuals or groups exhibit in a physical space that can be directly or indirectly measured to inform the control objective. The main challenge at this level is to create computational models that accurately simulate human behavior dynamics and decision processes for system design and control. Major applications include collaboration between humans and robots, rehabilitation and assistance, self-driving vehicles, process plants, smart homes, and managing the ambient environment.

However, in deriving computational models that describe the interaction between humans and CPS, there are no first-principle models to draw upon—this is in stark contrast to pure CPS. Hence, various simulation approaches have been utilized to analyze human behavior in the built environment, each requiring human behavior models at different levels of abstraction. As identified by Schaumann and Kapadia (2019), these approaches include system dynamics, process-driven, flow-based, particle-based, and multi-agent simulation (MAS). MAS is distinguished by its ability to provide detailed simulations that capture the complex interactions between agents, using specific behavioral rules that reflect their motives and preferences, enabling autonomous decision-making. This method facilitates the emergence of behaviors that are more representative of real-world complexities, offering a valuable tool for understanding and predicting human behavior in dynamic settings. Nonetheless, formulating predefined rules for agents to effectively engage with each other and their changing environments is challenging due to the range of factors influencing behavior and preferences, including physiological (e.g., age, 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 norms), social ties, and random factors Doctorarastoo et al. (2024b).

⋆ This material is based upon work supported by the National Science Foundation under Grant Number 2425121.

⋆⋆Mario Bergés holds concurrent appointments at Carnegie Mellon University (CMU) and as an Amazon Scholar. This manuscript describes work at CMU and is not associated with Amazon.

† Corresponding author: Katherine A. Flanigan

10.1016/j.ifacol.2025.01.166

Researchers have long tried to develop the rules governing spatial behavior based on these influential factors. These efforts have led to the development of different modeling approaches, predominantly classified as knowledge-driven (Chun et al. (2019)) and data-driven (Zhu et al. (2023)) approaches. Knowledge-driven models are traditionally built on data sourced from surveys, observational studies, and expert insights (Dawes and Ostwald (2014)). On the other hand, the rise of the Internet of Things and advancements in sensor technology have transformed data collection, enabling the assembly of detailed datasets that fuel data-driven models (e.g., privacy-preserving activity recognition) (Martins et al. (2023); Lin et al. (2024b)). The data used for training these models incorporate both known (e.g., temperature) and unknown variables that can affect decisions, yet these models often struggle to generalize to environments not included in their training data. Hybrid models have the potential to merge the advantages of each modeling approach, reducing their individual limitations and amplifying their strengths. In this context, Reinforcement Learning (RL) stands out as a primarily data-driven technique that also employs stochasticity. RL's strengths in complex and evolving environments stems from its dynamic learning capabilities and its proficiency in managing complexity and uncertainty, focusing on long-term goal optimization and the strategic balance between exploration and exploitation. RL also mirrors human decision-making processes by being goal-driven, sequential, and focused on maximizing the internal utility functions of agents.

However, in the context of modeling human behavior, this internal utility function, which reflects individual goals and preferences, is unknown. Therefore, the task of formulating humans' internal utility functions based on data remains a complex challenge. This is not only due to the individual-specific factors that influence human decisions, but also because the human agent must perceive relational configurations—that is, the spatial arrangement and relative positioning of elements and people within the environment with which the agent interacts. If we can model human perceptions and reactions to the built world spatio-temporally, we can apply control actions to these configurations to encourage specific behaviors, aligning with the objectives of CPHS.

To meet the multi-faceted needs of accounting for factors influencing human decisions *and* human agents' perception of relational configurations within their environment, we use Graph Neural Networks (GNN) in our study. GNN can uniquely capture, process, and aggregate relational data Scarselli et al. (2008); Taghizadeh et al. (2024), offering notable advantages in representing both the physical and social features of the environment and their interconnections. Therefore, we use GNN as a key building block to map complex human perceptions and preferences to reward functions. The proposed GNN-based preference model enhances realism of the reward function in the RL algorithm, compared to allocating rewards merely for performing activities. This results in behavioral simulations that are responsive to the diverse physical, social, and environmental attributes of the built environment.

In the remainder of this paper, we first explore the construction and training of the GNN-based preference model. Next, we investigate the use of RL to autonomously learn action rules for modeling human spatial behavior in dynamic settings. We then detail how this GNN-based preference model is integrated with the RL-driven behavior model to form a preference-aware human spatial behavior model. This multi-layered approach improves the accuracy of behavior simulations by incorporating individual preferences, while also enhancing adaptability through RL's ability to adjust to changing conditions. To showcase the practical applications and benefits of this model, we present an illustrative example featuring two scenarios in a conference room at Carnegie Mellon University (CMU) in Pittsburgh, PA. Using synthetic data, we demonstrate how this model improves simulation realism and behavioral predictions compared to a non-preference-based approach.

2. METHODOLOGY

2.1 GNN Preference Model

In this section, we systematically outline the development of the proposed preference model. First, in "Data Collection," we develop a multi-layered dataset that captures the spatial, social, and environmental factors within built environments. Then, in "Graph Representation of the Environment," we transform the gathered data into a graph structure in preparation for analysis with the GNN. Building on this graph-based representation, the "Implementation of the Preference Model" section discusses how the dataset is used to develop the preference model.

Data Collection. A key input of the GNN is *in-situ* data collected from sensors, such as from depth or RGB cameras. The authors note that in public settings (e.g., smart cities), there is an ongoing push to generate valuable data using privacy-preserving technologies, undergoing processing through activity recognition algorithms, as detailed in Martins et al. (2023) and Lin et al. (2024b). The data is further enhanced by adding tracklets capturing the exact locations of the activities observed. Additionally, a *social map* is created by tracking the location of individuals in the area over time. The environment's ambient conditions are monitored using sensors (e.g., digital thermometers, photodiodes), providing real-time environmental data. Data layers, such as time stamps recorded by system clocks and the physical layout documented through digital blueprints or LiDAR scans, are crucial for building a comprehensive dataset. This multifaceted dataset lays the groundwork for the GNN model to analyze and interpret human preferences within spatial and social contexts of the built environments.

Graph Representation of the Environment. To utilize GNN, the collected data in the "Data Collection" section must be represented in a graph structure. Given the spatial setting, the environment is segmented into a grid, where each grid cell is mapped to a node of a graph, $G = (V, E)$. Here, V represents the set of nodes, each corresponding to a distinct grid cell within the built environment. The set E comprises edges illustrating the proximity between the cells, with an edge, e_{ij}, existing in E if there is a direct adjacency between cells i and j. The graph's adjacency matrix, A, is defined such that $a_{ij} = 1$ if there is an edge $e_{ij} \in E$, indicating adjacency between the two cells, and $a_{ij} = 0$ otherwise. We denote the number of nodes and edges by $|V| = N$ and $|E| = L$, respectively, which depends on the refinement of gridification and the size of the environment. Each node $v_i \in V$ is associated with a composite feature vector, encapsulating cell i's spatial, environmental, and social characteristics, x_i:

$$x_i = [Er_i, En_i, S_i] \tag{1}$$

where Er_i is a binary vector indicating ergonomic features, with the binary format simplifying the representation of cells containing multiple items. Vector En_i reflects the built environment's environmental attribute, while S_i indicates whether the cell is currently occupied by other individuals. Vectors Er_i and En_i are adaptable to include a variety of features depending on specific applications. For example, in educational settings, Er_i may include elements like the presence of tables, chairs, and proximity to walls, windows, and exits, whereas En_i could incorporate factors like temperature, light intensity, noise level, and humidity.

In GNN, the term *target* refers to the output or ground truth of a node, edge, or the entire graph. This depends on the specific task, whether it is node classification, graph classification, or edge prediction. Our approach utilizes node classification within the GNN model to determine the probability of each grid cell (node) being selected for various activities. Consequently, each node is assigned a label indicating whether that cell has been chosen for a specific activity. The target vector for each node i is defined, where each element corresponds to one of the M possible activities:

$$\mathbf{y}_i^T = [y_{i1} \; y_{i2} \; \cdots \; y_{iM}] \tag{2}$$

Implementation of the Preference Model. The development of the preference model unfolds in two main phases: *Training* and *Prediction*. During the *Training* phase, the model assesses grid cells and calculates probabilities that reflect their likelihood of being selected for various activities. This evaluation uses preprocessed input and output data, as detailed in the "Graph Representation of the Environment" section. The preference model integrates each cell's spatial, environmental, and social features into a graph-structured format, referred to as the *Input Graph*. The *Output Graph* illustrates the selection of specific cells based on the *Input Graph*. This selection mechanism is shaped by the dynamics of human physiological and psychological traits and their preferences in terms of relational configuration of the physical space, captured through sensor data, mirroring individual preferences in selecting certain cells. The preference model operates as follows:

$$H^{(0)} = X \tag{3}$$

$$H^{(l+1)} = \mathrm{ReLU}\left(GNN^{(l)}\left(H^{(l)}, A\right)\right) \tag{4}$$

$$\hat{Y}_v = \sigma\left(GNN^{(L-1)}\left(H^{(L-1)}, A\right)\right) \tag{5}$$

where at each layer l, $H^{(l)}$ denotes the node features, and A represents the adjacency matrix that indicates the connections between nodes. The GNN layers, indexed by l, refine these features to calculate the probability $\hat{y}_v$ for each node's selection, utilizing the Sigmoid function σ. The GNN's parameters are then adjusted to ensure the model's predictions align with the target vector. Given the imbalance present in the training data, where only one node is selected from the entire graph, we employ 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. The loss is computed across all nodes N in the environment graph and averaged over the training samples T and the total number of nodes N as follows:

$$\begin{aligned} L_{BCEw}(Y_v, \hat{Y}_v) = -\frac{1}{T \cdot N} \sum_{j=1}^{T} \sum_{i=1}^{N} \Big[w \cdot y_j^{(i)} \log(\hat{y}_j^{(i)}) \\ + (1 - y_j^{(i)}) \log(1 - \hat{y}_j^{(i)}) \Big] \end{aligned} \tag{6}$$

where Y_v and $\hat{Y}_v$ denote the true labels and the estimated probabilities for nodes of the graph, respectively. The variable w represents the weight given to the positive class. After successfully training the GNN, the model moves into the *Prediction* phase. During this phase, it calculates the probability of selecting each grid cell for any of the M possible activities referred to as p_m.

2.2 Preference-Aware RL Spatial Behavior Model

Specification of RL Behavior Model. In RL, agents learn optimal behaviors through sequential decisions by interacting with their environment and receiving feedback in the form of rewards or penalties. The core components of RL include agents, the environment, actions, states, and rewards (Sutton and Barto (2018)). In our work, humans are represented as RL agents, with their spatial decision-making behavior modeled using policy functions. The environment encompasses the cyber, physical, and social worlds in which the agent operates and interacts. The human agent observes the state of the environment, which may include partial information from the surrounding physical, social, and cyber worlds, referred to as the observation. Based on this observation, the RL human agent performs specific activities through fine-grained actions such as moving, sitting, standing, and initiating interactions. The agent's objective is to successfully accomplish these activities, for which it receives rewards. However, such a simplistic reward function ignores humans' spatial preferences for performing activities, resulting in homogeneous behaviors across all agents that focus solely on the execution of activities. This does not reflect how humans behave, an issue addressed in the following section.

GNN Integration with RL. The integration process begins with the RL agent observing the environment's current state. Using the GNN's predictions on the likelihood of spaces being chosen for different activities, the agent evaluates the desirability of various locations rather than relying on a simplistic reward function. This means that the rewards received by the agent are determined by its complex perceptions and spatial preferences captured by the GNN. During each time step, the agent's policy function evaluates possible actions based on the GNN-predicted rewards indicating the suitability of locations, resulting in human agents choosing locations that align with their specific preferences. For instance, if the GNN predicts a high likelihood for a particular location being ideal for a specific activity (e.g., studying in a quiet, well-lit corner), the RL agent receives a higher reward for choos-

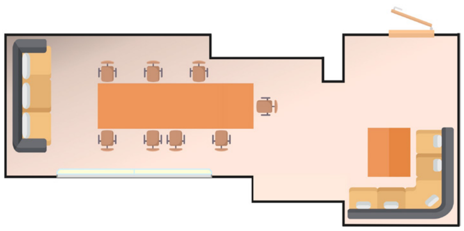

Fig. 1. Non-grid floor plan (Doctorarastoo et al. (2024a)).

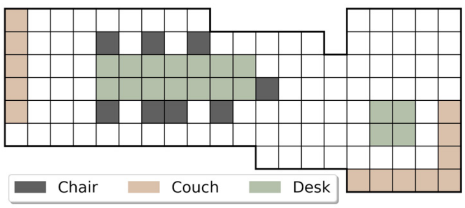

Fig. 2. Grid floor plan (Doctorarastoo et al. (2024a)).

Table 1. Light interpolation parameters.

Location	Mean (μ)	Standard deviation (σ)
Next to window	600 lux	200 lux
Corners	300 lux	150 lux
Under artificial lights	500 lux	75 lux

ing that location for studying. Conversely, less suitable locations result in lower rewards. This reward can change based on the evolving physical and social characteristics of the surroundings, potentially prompting the agent to change their location if suitability notably changes, guiding the agent toward more human-like spatial behaviors.

3. ILLUSTRATIVE EXAMPLE

A practical application of CPHS is within infrastructure systems, referred to as cyber-physical-social infrastructure systems (CPSIS). CPSIS extends traditional CPS by integrating the interactions between humans and infrastructure (Doctorarastoo et al. (2023a,b)). In CPSIS, physical infrastructure is designed or controlled to achieve not only economic goals but also human-centered objectives. For instance, how can we design and operate educational spaces to maximize productivity and learning quality? Or, how can we manage heating, ventilation, and air conditioning systems to minimize energy use while maximizing comfort? As an illustrative proof of concept, we demonstrate the capabilities of our proposed work within this context. This case study develops a human preference model in an educational facility at CMU, serving as a compact model for examining human interactions and preferences in a well-defined spatial environment. Within the broader control paradigm, this is an important aspect of modeling needed to meet objectives desired by the community such as productivity, collaboration, or sociability.

Case Study Setting. The conference room, shown in Figure 1, is a 4m×10m space that has been divided into a grid (Figure 2), converting the room into a discrete spatial model for graph representation. Each 50cm×50cm cell within the environment grid is assigned values representing physical entities like walls, entrances, windows, movable objects (e.g., chairs, desks), and unoccupied spaces. The identified activities are studying, eating, and socializing.

Table 2. Temp. interpolation parameters.

Location	Mean (μ)	Standard deviation (σ)
Next to window and heating	26°C	3°C
Corners	22°C	1.5°C
Next to entrance	19°C	2°C

Table 3. Preference model training parameters.

Parameter/method	Value/description
Optimizer	Adam
Learning rate (LR)	0.0005
Batch size	128
Activation function	ReLU (hidden layers), Sigmoid (output layer)
Loss function	Weighted BCE
LR scheduler	Step size 50 and gamma 0.99
Early stopping patience	50 epochs
Number of epochs	1000 (max)

Graph Structure. A subset *mesh* is extracted from the grid, encompassing all cells excluding those marked as unoccupiable or as walls, focusing on potential areas for human presence. A graph structure overlays the mesh, with nodes representing mesh cells and edges connecting adjacent nodes, including diagonals. Each node possesses a feature vector encompassing environmental characteristics (e.g., light intensity, temperature), object presence (e.g., tables, chairs), proximity to architectural elements (e.g., windows, walls), and occupancy by individuals.

Creating the Synthetic Dataset. For this proof-of-concept illustration, we generated 10,000 synthetic data points per activity to analyze spatial preferences. Node feature vectors are populated based on fixed room characteristics, while social and environmental attributes (referred to as maps) are derived using probabilistic distributions to better mirror real-world conditions and to enrich the dataset. The social map indicates individual locations, generated by randomly placing individuals to simulate presence of others, influencing preferences for social interactions or quiet zones conducive to concentrated tasks. Light intensity values are assigned to nodes based on their distance from light sources, i.e., windows and artificial lighting. The temperature map reflects room temperature distribution influenced by proximity to windows, heating, and entrance. Table 1 and Table 2 detail the Gaussian probability distribution parameters at specific locations for light intensity and temperature, respectively. These values are then interpolated to other locations using Nearest Neighbor interpolation for light intensity and Cubic interpolation for temperature.

Human preferences for activities are generated using a probabilistic selection mechanism that evaluates cells for suitability as activity zones. Suitability is quantified through a weighted sum of node features, with weights reflecting the relative importance of features for a specific activity, which can indicate willingness, indifference, or avoidance to different extents. For example, when studying, one might prioritize light intensity and proximity to windows, while for socializing, they might prioritize furniture arrangement. Each data instance involves selecting an initial node representing an activity's starting point and a target node chosen based on weighted feature preferences, encoding preferences given the current spatial, social, and environmental features. For this case study, we included three human agents in the conference room. *Agent 1*, while studying, prefers areas with couches, avoids high temper-

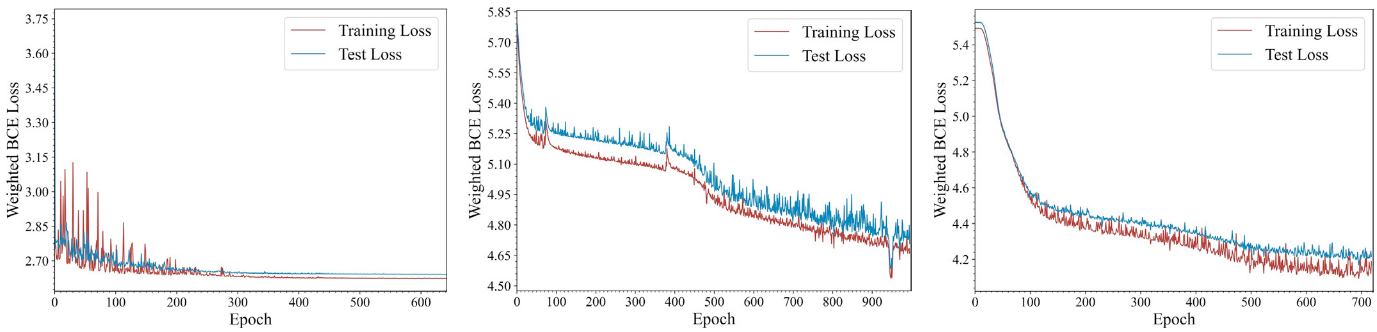

(a) (b) (c)

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

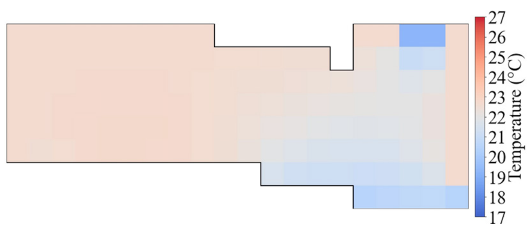

Fig. 4. Temperature heatmap of the conference room.

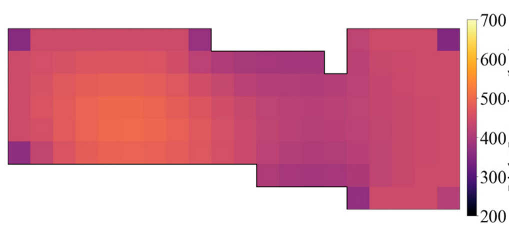

Fig. 5. Light intensity heatmap of the conference room.

atures, 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 by iteratively adjusting its parameters using the Adam optimizer to minimize the difference between the predicted outcomes and the actual data. Early stopping is applied to avoid overfitting. Table 3 summarizes key training parameters. The decline in training and test loss curves (Figure 3) indicates the model's effective learning and generalization. Their convergence suggests a balance between adapting to the training data and generalizing to unseen data.

4. RESULTS

In this section, we present the results of the proof-of-concept demonstration, showing movement trajectories and usage patterns for the three agents using simulation. The environment's temperature and light intensity maps are shown in Figure 4 and Figure 5, respectively. The analysis of agent trajectories in the conference room reveals distinct behaviors based on the inclusion or exclusion of preference modeling. Figures 6a, 6b, and 6c show the trajectories for the three agents across two scenarios: with and without preference modeling. Each scenario includes 200 trajectories to provide a robust comparison.

When preferences are not considered, the agents tend to select the nearest available seating to the entrance, leading to a higher concentration of activity in this area. This behavior aligns with a basic utilitarian approach where the primary goal is minimizing travel distance within the room. Consequently, the closest chairs to the entrance are often occupied first, reflecting a simplistic and predictable pattern of movement. Introducing the GNN-based preference model significantly alters the agents' behaviors, as agents seek out locations aligned with their individual preferences regarding temperature, light, and social distance. This leads agents to exhibit behaviors that more closely mirror human decision making, leading to diverse and realistic movements within the space. The results show that agents in the preference-aware scenario often traveled farther distances to reach their preferred locations. This finding suggests that preference-driven decision making can result in longer, but more goal-oriented, paths—an important consideration for optimizing environments, such as in urban planning or smart building systems.

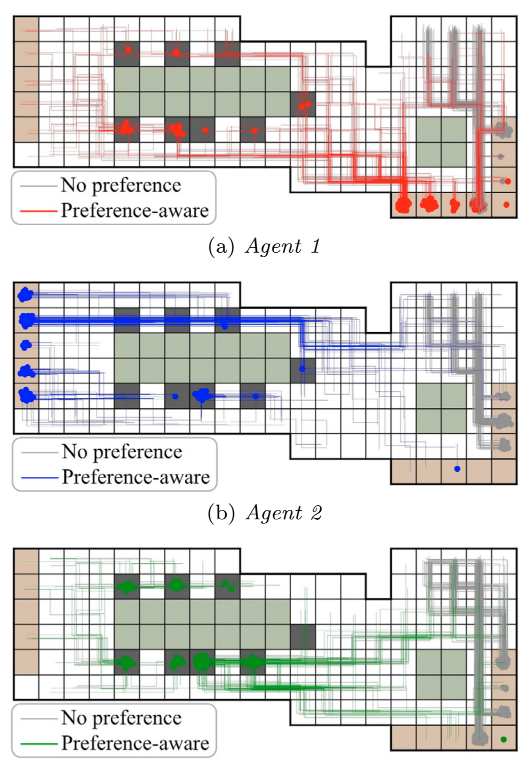

(a) *Agent 1*

(b) *Agent 2*

(c) *Agent 3*

Fig. 6. Agent movement trajectories showing behavior with and without preference modeling.

Agent 1, who prefers couches and cooler areas, predominantly chooses the couch in the cooler part of the room, prioritizing temperature comfort over proximity to the entrance. *Agent 2*, who favors solitude and warmer temperatures, moves across the room to the couch located in far left, clearly favoring both thermal comfort and social distancing. *Agent 2*'s decision is influenced by the presence of other agents, particularly *Agent 3*, reflecting the multi-agent nature of the simulation. As shown in Figure 6c, *Agent 3*, who prefers well-lit areas and proximity to windows, consistently selects seating near the window, favoring the chairs closer to the window over the more distant ones. This prompts *Agent 2* to select the couch, even though the temperature near the chairs is similar to the couch and they are positioned closer to the entrance. This interaction demonstrates how preferences for social distance can outweigh simple spatial proximity.

5. CONCLUSION

This paper introduces a new approach to modeling human spatial behavior in CPHS, with a focus on preference-aware decision making. By integrating GNN, RL, and *in-situ* sensor data, the model captures individual preferences related to environmental, social, and ergonomic factors, resulting in more realistic simulations of human behavior. This approach represents a significant advancement over traditional models by accounting for the dynamic nature of human-environment interactions and individual preferences. This multi-layered approach holds promise for revolutionizing areas such as smart building management, urban planning, and human-robot collaboration, where human-centric design and adaptability are crucial.

Future work should prioritize incorporating near real-time model retraining to adapt to evolving preferences and goals. Another key area for exploration is evaluating how well the training data represents human physiological and psychological traits, with comparisons to existing benchmarks. Expanding the model to account for internal factors such as emotions (Lin et al. (2024a)), refining the reward discount factor to balance short- and long-term objectives, and validating the models using real-world data will also improve this work's applicability.

REFERENCES

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

Chun, J., Fatah gen. Schieck, A., Psarras, S., and Koutsolampros, P. (2019). Agent based simulation for 'choice of seats': A study on the human space usage pattern. In *Proceedings of the 12th Space Syntax Symposium*, 1–18. International Space Syntax Symposium, Beijing, China.

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

Doctorarastoo, M., Flanigan, K., Bergés, M., and Mccomb, C. (2023a). 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*, 385–389. Istanbul, Turkey. doi:10.1145/3600100.3626340.

Doctorarastoo, M., Flanigan, K., Bergés, M., and McComb, C. (2023b). 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*, 370–376. Istanbul, Turkey. doi:10.1145/3600100.3626338.

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

Doctorarastoo, M., Lin, C., Flanigan, K., and Bergés, M. (2024b). 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*. Hangzhou, China. doi:10.1145/3671127.3699530.

Lin, C., Doctorarastoo, M., and Flanigan, K. (2024a). 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*. Hangzhou, China. doi:10.1145/3671127.3699529.

Lin, C., Martins, J., and Flanigan, K.A. (2024b). 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)*, ASCE. Pittsburgh, PA, USA.

Martins, J., Flanigan, K., and McComb, C. (2023). 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*, 377–384. Istanbul Turkey. doi:10.1145/3600100.3626339.

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

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

Schaumann, D. and Kapadia, M. (2019). Modeling social and spatial behavior in built environments: Current methods and future directions. *Social-Behavioral Modeling for Complex Systems*, 673–695.

Sutton, R.S. and Barto, A.G. (2018). *Reinforcement learning: An introduction*. MIT press.

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

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