Simulation and Prediction of Human Flows in Hospital Environments from Assistive Robot Observations

Maxime Bernard[1], Damien Rambeau[1], Jacques Saraydaryan[1,2], Olivier Simonin[1]

***Abstract*— Predicting human locations and flows is a key factor in social robotic navigation. This paper presents a framework for modeling and predicting individuals' motion patterns from local robot observations in hospital environments. We extend the human flow grid (Flowgrid) model to distinguish between agent categories (e.g., doctors, nurses, and patients) and propagate their presence probabilities using typed Flowgrids. Time-aware Flowgrid management further enables adaptation to the temporal rhythms of hospital operations. The framework is demonstrated in a simulated hospital environment with a multi-robot team.**

I. Introduction

Social robot navigation in human environments remains challenging due to uncertainty in human motion, social conventions, and dynamic crowd behavior [1]. Hospital environments are particularly demanding, combining narrow corridors, heterogeneous agent populations (staff, patients, visitors), and time-varying circulation patterns that purely reactive strategies cannot anticipate; long-term spatio-temporal maps of pedestrian flows have therefore been proposed to enable predictive navigation beyond the robot's sensing horizon [2].

Maps of Dynamics (MoDs) [3] encode place-dependent motion patterns from prior observations. Our prior work introduced the Flowgrid [4], a discrete directional flow map recording human motion probabilities per grid cell, extended with a gaussian-aided presence probability propagation model (HPP) [5] validated for social navigation.

While agent-type-aware social costmaps have been explored [6], the use of *typed flow maps* : distinct Flowgrids per agent category constraining presence propagation, has not been previously addressed. This paper introduces novel typed Flowgrids and time-aware map management within a multi-robot architecture for a hospital simulation with heterogeneous agent roles.

Sections II–IV describe the simulated hospital environment, Flowgrids, and the HPP model respectively; Section V presents experiments and Section VI concludes.

II. Pedestrian Simulation in a Hospital Environment

A. Hospital Environment

The experiments are conducted within the SOLAR-Nav project, which targets multi-robot navigation in a simulated hospital environment. This project is in collaboration with the "Hospices Civils de Lyon" which host a robotic experimentation platform that can provide maps of hospital services and ground-truth data on people's behaviour in these locations. The hospital map comprises typical clinical zones : corridors, patient wards, medical staff areas accessible to doctors and nurses only, waiting rooms, and shared circulation spaces. This spatial heterogeneity is central to the problem : circulation patterns differ substantially between zones, and certain areas are structurally restricted to specific agent categories.

B. Pedestrian Simulation with BePed

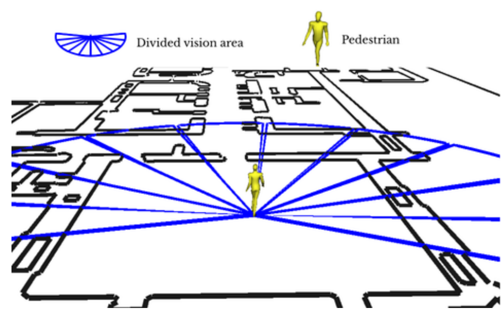

Fig. 1. Agent perception field divided into ten angular sectors; each sector independently computes a repulsive force from the agent's nearest obstacle.

Pedestrian are simulated using BePed, our extension to SPACiSS (Simulator for Pedestrians and an Autonomous Car in Shared Spaces) [7], itself built on Pedsim [8] and the Social Force Model (SFM) [9]. SPACiSS was originally designed for open-space crowd simulation and does not handle structured indoor environments well. BePed extends it to support corridor navigation and door crossing by modifying the obstacle force model: rather than reacting to the single closest wall, each agent's perception zone is divided into 10 angular sectors (see Fig. 1), and repulsive forces from the nearest wall in each sector are summed, producing smoother navigation through corridors and doorways.

Pedestrian navigation is also extended beyond simple waypoint-following: rather than applying a direct attractive force toward a goal which could trap the agent in rooms or corners, BePed computes planned paths using a Voronoi skeleton of the environment combined with A* search, producing trajectories that naturally follow corridors with respect to the building topology.

BePed introduces the definition of heterogeneous agent types, each associated with a behavioral profile, a set of goals, and a set of accessible zones. Three categories are modeled in the hospital scenario:

*This work was supported by the ANR agency France 2030 ANR-23-DMRO-0018 "SOLAR-Nav" project.

[1]Inria Lyon, INSA Lyon, CITI Lab., Chroma team, Villeurbanne, France.
[2]CPE Lyon, CITI Lab., Chroma team, Villeurbanne, France.

- **Doctors** circulate between consultation rooms, treatment areas, and staff zones following schedule-driven trajectories.
- **Nurses** exhibit broader circulation, regularly moving between patient wards, storage areas, and corridors.
- **Patients** move more slowly and are confined to patient-accessible zones such as wards and waiting areas.

In order to make the agents' behavior more realistic, they are governed by a behavior tree which manage their destination, their interactions with one another, and with objects, such as chairs or beds. This produces the varied, role-specific circulation patterns that motivate the use of typed Flowgrids (see Sect. III), as well as time-varying agent densities that reflect hospital operating rhythms: shift changes generate characteristic flow peaks in corridors, while off-peak periods produce sparser and more confined circulation.

C. Robot Platform and Observation Pipeline

To assist people, we deploy several Mirokai robots, autonomous ballbots capable of navigating the simulated hospital environment. Each robot is equipped with a detection pipeline that estimates the position and heading of observed agents, along with their category label. Observations are shared across the fleet via a centralized aggregator, feeding the HPP map update process described in Section IV.

III. Flowgrid: Mapping Human Circulation Patterns

A. Definition and construction

For any agent observed in a cell, we model the probability of transition toward each of the eight neighboring directions. The Flowgrid [4] is defined as a grid in which each cell stores a directional flow probability over the eight cardinal orientations: N, NE, E, SE, S, SW, W, NW.

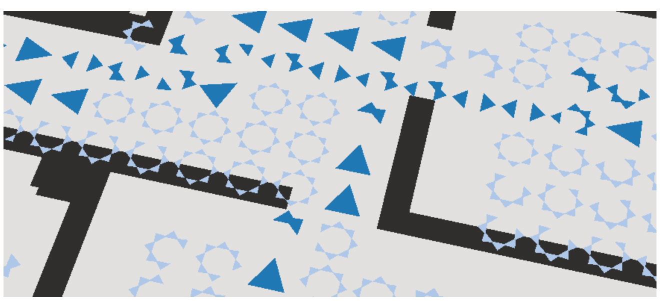

Fig. 2. Flowgrid representation. Each cell displays arrows (triangles) pointing the non-zero flow probabilities in each of the eight cardinal directions. A single arrow indicates all observed motion was in one direction; multiple smaller arrows reflect multi-directional flow. Lighter arrows denote directions extrapolated via the Von Mises model rather than directly observed (cf. Section IV, [5]).

Formally, let $R = \{r_1, \ldots, r_n\}$ be a set of robots observing the environment, $t_{c_{x,y}}(r)$ the total observation time of robot r on cell $c_{x,y}$, and $t_{c_{x,y,k}}(r)$ the total observation time in direction k on that cell. The Flowgrid value for cell $c_{x,y}$ in direction k, estimated from the observation history Z^t, is:

$$M_{c_{x,y,k}}(Z^t) = \frac{\sum_{n=1}^{|R|} t_{c_{x,y,k}}(r_n)}{\sum_{n=1}^{|R|} t_{c_{x,y}}(r_n)} \qquad (1)$$

The Flowgrid, illustrated in Fig. 2, is a static representation of circulation habits, built offline from recorded trajectory databases and loaded at runtime or updated online, used to guide HPP propagation (cf. Section IV). Flowgrids are stored as 3D matrices where the first two dimensions encode spatial position and the third encodes orientations.

B. Typed Flowgrids per agent category

In hospital environments, a single global Flowgrid combines the behaviors of all agent types, producing an averaged representation that can mislead presence estimation. We associate each agent category with a dedicated *typed Flowgrid*, built from trajectories filtered by type and illustrated in Fig.3.

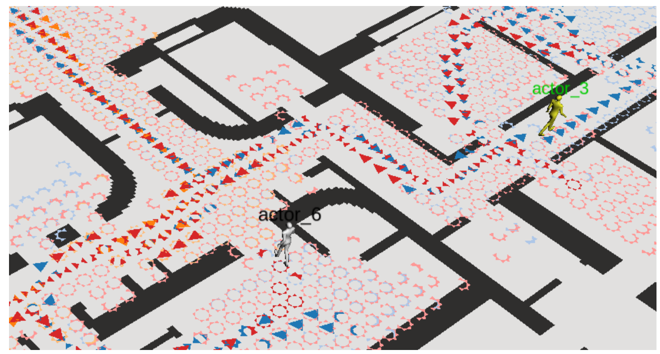

Fig. 3. Typed Flowgrids per agent category (doctors, patients...), each built from trajectories filtered by agent type. Colors distinguish agent categories.

A global flow map can be recovered by summing per-orientation probabilities across typed Flowgrids, but the key novelty lies in using each typed Flowgrid independently during HPP propagation: a detection of a patient, for instance, triggers propagation guided by the patient Flowgrid only, ensuring that presence probability does not flow into zones structurally inaccessible to that category (eg. staff-restricted areas).

C. Time-Aware Flowgrid Management

Hospital flows vary significantly throughout the day: shift changes, consultation schedules, and visiting hours produce characteristic temporal signatures in each zone. To capture these patterns, typed Flowgrids are constructed from SQL-parameterized queries on trajectory databases, filtering observations by time slot or day of the week. Precomputed Flowgrids associated with specific time windows are saved and loaded at runtime, enabling fast switching between flow maps as the operating context evolves.

IV. HPP Map: Real-Time Human Presence Estimation

A. Propagation model

The Human Presence Probability (HPP) map complements the static Flowgrid with a dynamic shared representation to predict human presence. Like the Flowgrid, the environment is discretized into cells; each cell holds a scalar presence probability, updated in real time according to the propagation model introduced in [5].

Mobile robots detect agents and share observations across the fleet. Given a discrete observation $z^t_{x_0,y_0,k_\theta}$ of an agent on cell c_{x_0,y_0} with orientation k_θ, we denote by $P^{t+1}_{\text{pres}}\left(c_{x_1,y_1,k} \mid z^t_{x_0,y_0,k_\theta}\right)$ the presence probability on neighboring cell c_{x_1,y_1} in direction k at time $t+1$, with initial condition:

$$P^t_{\text{pres}}\left(c_{x_0,y_0,k_c} \mid z^t_{x_0,y_0,k_\theta}\right) = 1 \quad \text{if } k_c = k_\theta, \text{ else } 0$$

When propagating using the Flowgrid alone, probability is redistributed according to recorded flow proportions from eq.1, attenuated by a decay factor $\alpha \in [0,1]$:

$$P^{t+1}_{\text{pres}}\left(c_{x_1,y_1,k} \mid z^t_{x_0,y_0,k_\theta}\right) = \frac{M_{c_{x_1,y_1,k}}}{\sum_{i=1}^{\|K\|} M_{c_{x_1,y_1,k_i}}} \times \alpha P^t_{\text{pres}}\left(c_{x_0,y_0,k_c} \mid z^t_{x_0,y_0,k_\theta}\right) \tag{2}$$

The Von Mises distribution $\text{VonMises}(\theta_k, \theta_{k_c})$ provides a smooth circular probability over directions centered on θ_{k_c}, used here to weight propagation toward directions consistent with the observed orientation even in cells not yet covered by the Flowgrid. The combined model introduced in prior works [5], leveraging both the Flowgrid and the Von Mises orientation model, prioritizes propagation along observed flows while preserving non-zero probability in directions not yet recorded:

$$P^{t+1}_{\text{pres}}\left(c_{x_1,y_1,k} \mid z^t_{x_0,y_0,k_\theta}\right) = \frac{\text{VonMises}(\theta_k, \theta_{k_c}) \cdot M_{c_{x_1,y_1,k}}}{\sum_{i=1}^{\|K\|} M_{c_{x,y,k_i}}} \times \alpha P^t_{\text{pres}}\left(c_{x_0,y_0,k_c} \mid z^t_{x_0,y_0,k_\theta}\right) \tag{3}$$

where k_c is the orientation pointing toward cell c_{x_1,y_1}.

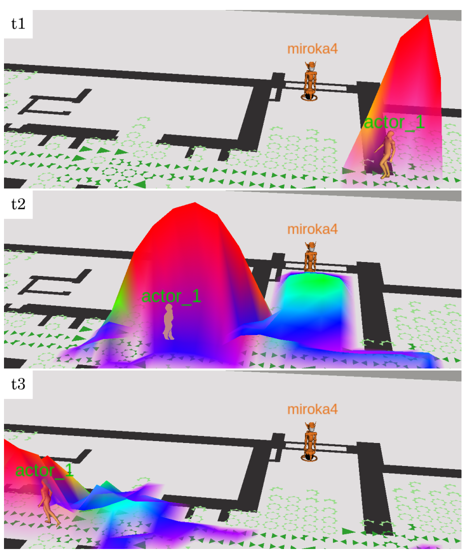

Fig. 4. HPP map propagation.

Fig. 4 illustrates the probability propagation with three keyframes; probability waves are rendered using a modified version of the *Grid Map Library* [10]. At t_1, the probability is initialized to its maximum value at the detection cell. At t_2 and t_3, it propagates primarily along the flows encoded in the Flowgrid, while also spreading into non-observed directions through the Von Mises extrapolation. This behavior is further illustrated in the video: https://youtu.be/-BST7qYwvA4.

B. Typed propagation

The integration of typed Flowgrids into the propagation model is the core contribution of the HPP extension presented in this work. When a robot observes an agent of a known category, propagation uses the Flowgrid specific to that category rather than a global map. This constrains presence estimates to the spatial patterns actually associated with the observed agent type, preventing cross-category artifacts: a nurse detection will not generate presence probability in patient-only areas, and vice versa. The HPP map is updated continuously by the fleet and shared across all robots, providing a human presence prediction throughout the environment.

V. Global Architecture and Experiments

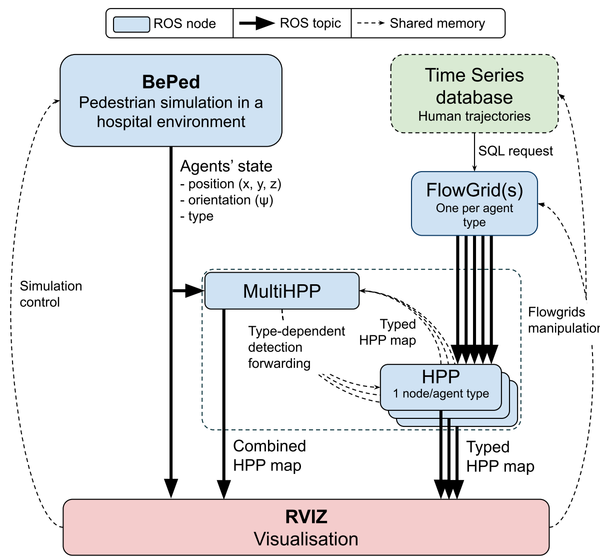

Fig. 5. System Architecture Diagram.

A. Global Architecture

Fig. 5 presents the functional diagram of the general architecture, illustrating its components and their interactions. A dedicated ROS 2 node handles the construction of typed Flowgrids, either online or via SQL queries on a human trajectory database. The resulting Flowgrids are forwarded to a set of presence probability propagation nodes, one per agent category, which receive agent detections from the BePed pedestrian simulator as input. An orchestrator node then aggregates the per-category HPP maps into a single global presence probability map. The full pipeline is visualized through an RViz scene extended with custom plugins providing control over the simulation state and time-slot management as described in Sect. III.

B. Hospital-Grounded Scenarios

The scenarios presented in this section were designed in collaboration with staff from the HCL (Hospices Civils de Lyon) hospital, whose field observations informed the definition of agent categories, their associated circulation zones, and the temporal rhythms of activity modeled in

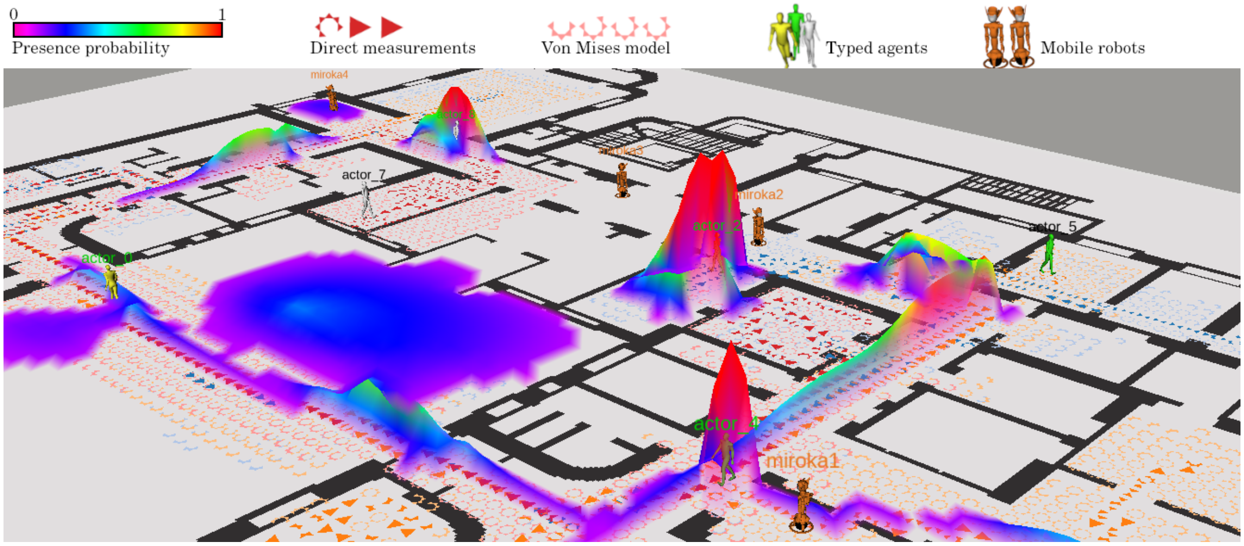

Fig. 6. Presence Probability propagation along typed Flowgrids, aided by the Von Mises orientation model.

BePed. Maps used in the simulator were built based on real hospital plans provided by the HCL. The experiments reported here provide preliminary validation of the core principles and models introduced in this paper, rather than a full quantitative benchmark.

C. Preliminary Results and Discussion

As shown in Fig. 6 and in the video available at https://youtu.be/-BST7qYwvA4, preliminary results illustrate the core mechanisms of typed Flowgrid construction and HPP propagation in a hospital-grounded simulation build on real field data. Agent trajectories were recorded from simulation runs and are used for the Flowgrid construction from database queries.

Each agent type is associated with a Flowgrid and an HPP map. While 3D matrices represent a consequent impact on the memory, our partners at HCL identified a limited number of agents groups, which, along the fixed hospital environment, ensures a linear scaling pattern when adding new agents category.

Simulated robots detection feeds the HPP propagation. Fig. 6 illustrates this behavior: recent observations generate high-probability regions shown in red to green, while older propagated estimates fade to lower-probability waves in purple to blue, reflecting the temporal decay modeled by α. From 02:09 in the video, the presence probability closely tracks the real agent's position while anticipating potential direction changes.

VI. Conclusion

This paper presented a framework for modeling and predicting human presence in hospital environments using a multi-robot team, grounded in real hospital data from the HCL and a behaviorally realistic pedestrian simulator.

The Flowgrid model is extended with two contributions: typed Flowgrids, one per agent category, which constrain presence propagation to category-specific circulation patterns and avoid cross-type artifacts; and time-aware map management, allowing the robot fleet to adapt to the temporal rhythms of hospital operations. Combined with the Von Mises-aided HPP propagation model, the framework provides a continuously updated, category-resolved estimate of human distribution shared across the fleet, directly addressing the pedestrian behaviour prediction challenges.

As future work, we plan to generate temporal cost maps for trajectory planning, leveraging HPP maps to estimate path costs as a function of predicted human presence, enabling robots to select trajectories that minimize interactions. We also intend to exploit typed Flowgrids for active agent search, guiding robots toward areas where the presence of a specific profile (doctor or nurse) is most probable. Finally, we aim to move toward hybrid simulation, connecting BePed to physical robots or cameras deployed on-site, enabling the framework to be updated from real observations.

References

[1] T. Kruse, A. K. Pandey *et al.*, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.

[2] T. Vintr, J. Blaha *et al.*, "Toward benchmarking of long-term spatio-temporal maps of pedestrian flows for human-aware navigation," *Frontiers in Robotics and AI*, vol. 9, p. 890013, 2022.

[3] T. P. Kucner, A. J. Lilienthal *et al.*, *Probabilistic mapping of spatial motion patterns for mobile robots*. Springer, 2020.

[4] F. Jumel, J. Saraydaryan, and O. Simonin, "Mapping likelihood of encountering humans: application to path planning in crowded environment," in *8th Euro. Conf. Mobile Robot, ECMR*. IEEE, 2017.

[5] J. Saraydaryan, F. Jumel, and O. Simonin, "Human presence probability map (hpp): A probability propagation based on human flow grid," in *Robot World Cup*. Springer, 2023.

[6] S. Kang, S. Yang *et al.*, "Social type-aware navigation framework for mobile robots in human-shared environments," *Sensors*, vol. 24, no. 15, p. 4862, 2024.

[7] M. Prédhumeau, "Simulating realistic pedestrian behaviors in the context of autonomous vehicles in shared spaces," in *AAMAS*, 2021.

[8] B. Okal, T. Linder *et al.*, "pedsim_ros: Pedestrian simulator powered by the social force model." GitHub, 2014. [Online]. Available: https://github.com/srl-freiburg/pedsim_ros

[9] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.

[10] P. Fankhauser and M. Hutter, "A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation," in *ROS – The Complete Reference (Volume 1)*, A. Koubaa, Ed. Springer, 2016. [Online]. Available: http://www.springer.com/de/book/9783319260525