Go with the Flow: Exploration and Mapping of Pedestrian Flow Patterns from Partial Observations

Sergi Molina, Grzegorz Cielniak and Tom Duckett

Abstract-Understanding how people are likely to behave in an environment is a key requirement for efficient and safe robot navigation. However, mobile platforms are subject to spatial and temporal constraints, meaning that only partial observations of human activities are typically available to a robot, while the activity patterns of people in a given environment may also change at different times. To address these issues we present as the main contribution an exploration strategy for acquiring models of pedestrian flows, which decides not only the locations to explore but also the times when to explore them. The approach is driven by the uncertainty from multiple Poisson processes built from past observations. The approach is evaluated using two long-term pedestrian datasets, comparing its performance against uninformed exploration strategies. The results show that when using the uncertainty in the exploration policy, model accuracy increases, enabling faster learning of human motion patterns.

I. INTRODUCTION

As robots increase their presence in environments shared with humans, such as care homes, hotels, warehouses or shopping centres, new challenges arise. It is important to know how the people are likely to behave and move in the environment so the robot can act accordingly, navigating in a more natural way. For example, a robot may use this knowledge to look for people if it has to interact with them, or go with the main direction of flow in a given area to minimise the possibility of collisions and increase safety.

Traditional exploration strategies assume a static environment. Their aim is to build a spatial model that covers the operational environment as quickly as possible, considering that a robot cannot see the entire environment at once. However, when modelling human motion, the environment tends to change over time, e.g. with different flow patterns at different times of the day. But, as in the spatial domain, a mobile robot cannot explore in a continuous manner due to temporal limitations, e.g. needing to recharge its batteries or perform other tasks not related to exploration. Hence, apart from deciding where to explore, it is also important to consider the temporal dimension (i.e. when to explore).

To address this problem, we present, as the main contribution, a mobile robotic exploration method which decides not only the locations from which to gather information about the people's motion but also the times of observations. The data captured by the robot is used to determine where and when to perform future observations based on the uncertainty levels of multiple time-varying Poisson distributions, which are built by counting detections of pedestrian motions at the chosen locations and times.

The main aim of this work was to build and refine a timedependent probabilistic map that is able to model and predict flow patterns of people in indoor environments. This model is also based on counting detections of pedestrians together with their observed directions of travel. For the experimental evaluation, we used two long-term term pedestrian datasets, which provide complex time-varying human motion patterns, to measure the performance of the novel exploration approach compared to other uninformed policies. One dataset was obtained from a shopping centre [1] and the other was recorded by the authors in office corridors at the University of Lincoln, providing 52 and 14 days of data, respectively. The results show that the proposed exploration system can learn the flow patterns and increase model quality more quickly than the corresponding uninformed strategies.

II. RELATED WORK

Exploration methods are aimed at guaranteeing model completeness by giving the robot ability to autonomously create an accurate model of the environment. One of the earliest approaches is frontier-based exploration [2][3], which drives the exploration towards the boundaries of the known and unknown until the map is complete. Another approach is adopted by next-best-view strategies, which focus on building the initial map based on optimising different criteria. For example, [4] optimises the estimated time to reach a location and the amount of information expected to be gathered there, [5] calculates the entropy decrease in the robot configuration space and uses the estimates for robot mapping, [6], [7] use the Poisson uncertainty of neutron counts to drive the observation and mapping of soil moisture, and [8] presented an information-gain-based exploration approach that takes into account the uncertainty from both the map and the robot's localisation. However, all the above approaches do not attempt to maintain the environment models after their acquisition, meaning that the model will lose accuracy as new changes appear in the environment.

To deal with this problem, another branch of algorithms aims to create models of the environment which allow them to predict where and when to make observations of specific phenomena, reasoning about the best times and locations. In [9] and [10], for example, the decision making for environmental surveillance and monitoring is based on Gaussian Processes, which allow the robot to learn the temporal patterns in the environment. Other approaches [11], [12] are instead based on the assumption that some of the environmental variations observed are caused by people's daily routines; [11] present a method for life-long spatiotemporal exploration of dynamic environments, using the entropy of binary state predictions in an occupancy map to create a scheduler that determines which areas and times to explore for each day. However, reasoning about human motion requires more complex model representations, and binary states are not sufficient to describe the pedestrian flows covered in this paper.

Modelling human motion has also been a recent topic of study; [13] present a probabilistic directional grid map to represent long-term angular motion, and [14] also takes in to account the observed speed of people. Despite being focused on representing human motion, these methods do no take into account temporal variations, instead modelling the average flow patterns, resulting in many cases in less accurate predictions. Thus, in the work developed by [15], periodic Poisson processes are used to characterise the behaviour in time of different rooms of a building.

Our exploration method is built on the concept presented in [16], which creates a time-dependent probabilistic map that is able to model and predict patterns of people in indoor environments. However, in [16], all the models were built assuming that the environment is fully observable both in time and space, which is unlikely in a real-world scenario.

III. SPATIO-TEMPORAL MODEL

In this work, we propose to use a representation that models the likelihood of motion directions on a grid-based map by a set of harmonic functions, which capture longterm changes of crowd movements over time. The model is updated after some time, e.g., end of every day, with the information gathered from the environment, which is then used to calculate the places and times to visit the next day.

A. Spatial representation

The underlying geometric space is represented by a grid, where each cell contains k temporal models, corresponding to k discretized orientations of people motion through the given cell over time. Since the total number of temporal models, which are of a fixed size, is $k \times n$ where n is the total number of cells, the spatio-temporal model does not grow over time regardless of the duration of data collection.

B. Temporal Framework - FreMEn

The temporal models, which can capture patterns of people movement, are based on the FreMEn framework [17]. Fre-MEn is a mathematical tool based on the Fourier Transform, which considers the probability of a given state as a function of time and represents it by a combination of harmonic components. The model not only allows representation of environment dynamics, but also prediction of future environment states based on the patterns learned. The idea is to treat a measured state as a signal, decompose it by means of the Fourier Transform, and obtain a frequency spectrum with the corresponding amplitudes, frequencies and phase shifts. Then, transferring the most prominent spectral components to the time domain provides an analytic expression representing the likelihood of that state at a given time in the past or future. Assuming that the directions of people movement are affected by patterns that might be periodic,

we apply the FreMEn concept to discretised directions of people movement through a particular cell.

C. Building the model

Our model assumes that it is provided with people detection data, comprising the person position, orientation and timestamp of the detection (x, y, α, t) . We associate each cell with k bins, corresponding to the discretised orientations of people motion, each with an associated (FreMEn) temporal model. When building the model, the x, y positions are discretised and assigned to the corresponding cell and the orientation α is assigned to one of the k bins, whose value is incremented by 1. After a predefined interval of time, we normalise the bins and use their values to update the spectra of the temporal models using the scheme described in [17]. Then, we reset the bin values to 0 and start the counting again. The same count values are also used to calculate the uncertainty measures, as explained further in the paper.

D. Making predictions

To predict the behaviour of human movement through a given cell at a future time t, we calculate the likelihood for each discretised orientation associated with that cell as

$$p_{\theta}(t) = p_0 + \sum_{j=1}^{m} p_j \cos(\omega_j t + \varphi_j), \qquad (1)$$

where p_0 is the stationary probability, m is the number of the most prominent spectral components, and p_j , ω_j and φ_j are their amplitudes, periods and phases. The spectral components ω_j are drawn from a set of ω_s that covers periodicities ranging from 1 to 24 hours with the following distribution: $\omega_s = s \cdot 3600$, $s \in 1, 2, 3, ..., 24$.

IV. EXPLORATION

We define an exploration strategy as the process that determines both the locations to visit and the times to do so, considering that a robot is limited by time and visibility constraints. The strategy has to deal with partial observability arising from the fact that a mobile platform only has a limited field of view due to sensor limitations. Also a mobile platform usually has to perform other tasks besides the environment exploration activities, so only a percentage of the total time, known as the exploration ratio, can be devoted to exploration. Thus, we need to define a metric that drives the exploration towards the areas that are more uncertain.

A. Defining uncertainty

When working with people, the appropriate probabilistic model for counting and modelling their occurrences is the Poisson distribution. Since the events occur independently, the number of events occurring in an interval can take only natural number values and two events cannot occur at exactly the same instant [15]. Hence, we propose to use the relative standard uncertainty from a Poisson distribution as an uncertainty measure, using this metric to define how unsure we are about the estimated people flows in a particular cell of the map during a given interval of time. We assume that all n cells in the map follow a Poisson distribution. So the probability of people detections occurring at every interval in every cell is distributed as:

$$P(j;\lambda) = \frac{\lambda^{j} e^{-j}}{j!}, \quad j \in 0, 1, 2, 3, ...,$$
(2)

where λ is the average number of detection counts in the corresponding interval, and j is the number of event occurrences.

The value λ for every cell evolves over time and so does Eq. (2), so in our case we work with inhomogeneous Poisson processes. As we want to work with partially observable environments, resulting in relatively few samples in each cell, we used a confidence-sensitive estimate for λ . In Bayesian inference, the conjugate prior for the rate parameter λ of the Poisson distribution is the Gamma distribution [18]: $\lambda \sim Gamma(\alpha, \beta)$, where α and β are the shape and inverse scale parameter, respectively. As we update λ for every interval, we can define the relative standard uncertainty of each cell as:

$$u_r = \sqrt{\lambda}/\lambda. \tag{3}$$

This measure infers that the more people we see within the boundaries of a cell, the more sure we become about how the people behave. An exception to this rule is introduced to define the uncertainty when $\lambda = 0$. According to the function, the uncertainty here would be infinite, however we introduce a threshold value to indicate that even though there is no one in that cell, in some way we are increasing the certainty of knowing that there is no one there. This threshold value (u_0) is chosen experimentally.

B. Exploration strategies

In this work, we test a total of 4 different exploration policies. The **Monte Carlo** strategy, which represents the main contribution of this paper, is driven by the uncertainty levels. This takes into account the experience from previously gathered data to infer the exploration actions at future times.

- In the Monte Carlo (MC) strategy, the places/times are chosen randomly, but the probability of selecting a given location/time interval is proportional to the uncertainty.

For comparison purposes we also implemented 3 different uninformed exploration strategies. In these cases, we do not take into account the environment dynamics, neglecting the uncertainty levels.

- The **Random** (**R**) strategy chooses the locations and times for exploration in an uniformly random way.
- The Round Robin (RR) strategy visits all the regions/time intervals of the environment with the same frequency.
- The Range (RA) strategy picks a region with probability proportional to the number of cells that can be seen inside.

In initial comparisons, we also implemented a **Greedy** exploration strategy, driving the exploration towards the region with the highest uncertainty. However, in all the experiments this method performed poorly compared to the other strategies. For this reason, we decided to omit the greedy approach from the results presented.

V. EVALUATION

A. Datasets

To evaluate the approach, we ran experiments using two real pedestrian datasets. Both feature complex human movement and enough days to train the models and evaluate the different exploration strategies in the long term.

The first one is a pedestrian tracking dataset recorded at the **ATC** shopping centre in Osaka, Japan [1]. From this dataset, we selected the first 52 consecutive days (26 Wednesdays and 26 Sundays), using 46 to perform exploration and the other 6 days as the evaluation data for result comparison. The recording of each day provides people trajectories starting from approximately 09:00 until 21:00, so for the rest of day we assume there are no occurrences of people.

The second dataset was collected at one of corridors in the Isaac Newton Building building at the University of Lincoln (**UoL**). The data collection was performed by a stationary robot equipped with a 3D lidar and a 2D laser. The robot was placed in a junction, covering a total area of around 75 m² (Fig. 1). To detect and track the people, we used an efficient and reliable method developed by Yan et al. [19]. Our dataset spanned from mornings to evenings for 14 days, sparsely recorded over a four week period. From these, 12 days were used for training and the remaining 2 for testing.



Fig. 1: UoL dataset: Robot location in the corridor and example of a person walking seen by the 3D lidar scans.

B. Model parameters

In our experiments, we discretised the space into 1×1 m regions for the ATC dataset and 0.5×0.5 m for the UoL one, resulting in a total of n = 1248 and n = 465 active cells, respectively. The number of bins chosen to discretise the orientations is in both cases k = 8, as shown in Fig. 2 in the upper right corner.

Moreover, we assume in our experiments that it is not possible to observe the state of the whole environment at every point in time, and instead we have defined a set of observable regions for each dataset. Each of the regions comprises a different number of reachable cells (shown in black in Fig. 2). For the ATC environment we defined 49 regions, each one measuring 7×7 m (in red in Fig. 2), while the UoL scenario comprises 10 regions of 5×5 m. During exploration only the people passing within the boundaries of the current region will be taken into account to update the model at a given point in time, and the rest of the environment is unseen.

Regarding time, we used the same parameters for both datasets. The interval for creating the histograms used as the



Fig. 2: Example of the shopping center map with 1 m cell grid and 49 regions (7×7 m) available for exploration. Upper right corner: 8 bins discretised at 45 degree intervals.

input for model creation was set to 10 minutes. The same interval was used to provide a single observation, i.e. every 10 minutes the policy can decide whether the robot should stay in the same region or instead move to a different region.

To build the Poisson processes in each cell we impose a global period of 1 day. This value comes from our previous work [16], which considered the same two datasets, where it was found that a daily rhythmic pattern gave the best fit. These periodicities were provided by the FreMEn tool, which analyses the temporal patterns of the data provided to it.

So in our case, the average people detection ratio is built using the same interval of time, for example from 10:00 to 10:10, across different days. As we use the conjugate distribution, the rate λ for each cell and time interval is updated by updating the Gamma distribution.

In these simulations using real pedestrian data we are currently not reproducing a robot moving from point to point, and it is assumed that the robot can observe the contents of only one of the regions for spatial exploration purposes. A more realistic robot simulation is beyond the scope of this paper, since we are mainly concerned with obtaining consistent results that enable direct comparisons of alternative exploration strategies.

C. Experiments

We divided the experiments into 3 parts. The first focuses on the problem of **where** to explore, the second on the problem of **when** to explore, and the last part considers **both** aspects.

In the **where** experiments we evaluate the strategies assuming that the space is partially observable, but the time is not. The mobile platform can dedicate all of its time to performing exploration (exploration ratio = 100%), however in each 10 minute interval it is possible to observe only one of the regions in the map. In the case of the Monte Carlo strategy, the region to explore is chosen based on the total uncertainty of each region, which is the sum of the uncertainties of all cells belonging to that region:

$$u_{region} = \sum_{c=1}^{n_{region}} \frac{\sqrt{\lambda_c}}{\lambda_c}.$$
 (4)

In the second part, we simulate the opposite situation, where the whole map is fully observable but we need to decide **when** to explore it based on time constraints. For



Fig. 3: Model predicition over 24h with m = 1, and probability distribution of each orientation at t = 18:00.

example, considering that a single day has 144 10 minutes intervals, setting the exploration ratio to 50%, means that 72 out of 144 will be used to update the models and the rest will be unseen.

Finally, we combine both cases and decide **both** where and when to perform the exploration. In all experiments, after the end of each day, the models and the uncertainties for the entire map are recalculated with the new measured detections.

VI. EXPERIMENTAL RESULTS

To evaluate the exploration strategies, we compare the predictions made by the model in the testing days against the ground truth at the end of each day of exploration. The output of the trained FreMEn models provides a function for each of the 8 orientations in every cell. So for each time t, we are able to obtain a distribution describing how likely it is to find a person moving in each direction. For example, Fig. 3 presents the prediction graph for a single cell in the map (using m = 1 FreMEn component) over 24 hours after some days of training. In this cell, there are 2 predominant orientations, one with higher probabilities in the morning and the other more prominent in the afternoon. At t = 18:00, the distribution of the 8 normalised orientation probabilities is shown in the polar histogram on the right, which is what we compare against the ground truth. However, obtaining the same orientation distribution with real data at a single time instance t is not possible, because we cannot count sufficient detections to build a meaningful distribution. Instead, the idea is to compare the distribution obtained with the predictions from the ground truth during a defined interval of time, which we set to 10 minutes in all the experiments.

To make the comparison between the predicted and actual histograms for each interval, we use the Chi-square distance to indicate the level of similarity. The higher the distance the less accurate is our model prediction compared to the ground truth. So at the end of every day during the exploration experiments, we compute the total error of all 144 intervals in all the corresponding testing days. The error of the whole map for a single interval is defined as:

$$e_{map} = \sum_{i=1}^{n} \sum_{b=1}^{k} \frac{(x_b - y_b)^2}{(x_b + y_b)},$$
(5)

where *n* is the number of cells, k = 8 is the number of angular bins for the direction of people motion in the cells, x_b is the value of bin *b* in the predicted orientation histogram, and y_b is the value of the same bin *b* obtained from the ground truth.

For all the experiments in this work, the predictions for the model are calculated using only 1 spectral component, as we impose only one periodicity.

The parameter u_0 was adjusted depending on the environment. For the ATC dataset it was defined as 1 during the 'where to explore' experiments and 0.1 for the 'when' case. For the UoL dataset, the parameter was set as 0.5 and 0.1, respectively. These values are specific to the environments and for the experiments carried out in this work.

In all the following plots, we provide a baseline model error corresponding to a fully observable environment, both in time and space (in black in Figures 4, 5 and 6). Also, since the MC, R and RA strategies are based on stochastic policies, the result is shown as the average together with a 95% confidence interval, using a Student's t-distribution, over 5 runs.

A. Where?

Fig. 4 shows the evolution of the errors in the model predictions using the different strategies in both environments. As a general overview, the results are consistent between datasets, with MC providing the best results and R the worst. However, all policies end up with similar accuracy values after the 46 and 12 days explored, respectively.

Observing the results of the ATC experiments, the MC and RA strategies perform closely in the first days, having a very similar learning ratio. By contrast, the R and RR strategies are slower in the first stage. Looking after day 10, we see that the RA strategy starts to fall behind the other methods, while MC maintains a slightly lower error. So biasing the exploration towards the regions where a larger number of cells can be seen is not always the recommended behaviour, as other regions, even with fewer cells, can be less certain.

For the UoL dataset, we notice that there is not as big a difference in the initial stage as in the previous dataset, however after day 5 MC maintains a lower error as before.

B. When?

The results shown in Fig. 5 correspond to experiments performed with 10% and 50% exploration ratios, so 14 and 72 intervals out of 144 are explored in each day, respectively.

In the 10% case, the MC and RR strategies are the ones performing better. In the ATC dataset, the best overall strategy is RR, mainly because in the first stage, the learning ratio is faster than the MC and R policies. We believe the reason behind the better performance of RR could be due to the temporal model being based on periodicities. With this policy, the robot explores the time domain in a structured manner, which helps to find the rhythmic patterns of the environment more quickly. With the UoL dataset, all three strategies perform similarly over the initial days, probably because the temporal rhythmic patterns are not as strong as

in the previous dataset. Only the R strategy seems to drop in performance after the 4th day with respect to the other strategies. For the 50% case, all strategies learn really quickly in both environments, and after just 5 days, the performance is very close to the baseline for all three strategies. The fact that we assumed full spatial observability in this experiment makes the choice of strategy almost irrelevant with high exploration ratios.

The RA strategy was not included as the number of observable cells would remain invariant across all the intervals. Since all the intervals would have the same probability, this strategy would behave the same as R.

C. Where and When?

For the last part of the experiments, the strategy decides both where and when to explore. The comparison is done with 5 different combinations, R-R, RR-RR, MC-MC, MC-RR and RR-MC, with the first named strategy indicating the policy for choosing the location and the second indicating the one in charge of deciding when to explore. Fig. 6 gathers the model prediction error over 46 days for the spatial and temporal combinations mentioned, for both exploration ratios (again 10% and 50%) and the two datasets.

As a general outline, independently of the dataset and exploration ratios, the strategies that include the uncertaintydriven exploration for the spatial domain (MC-MC and MC-RR) obtain a consistently better performance over their uninformed counterparts (R-R and RR-RR). This is more noticeable in the ATC environment with faster learning speeds, but in the UoL environment, MC-MC also ends with the lowest error among all strategies for both exploration ratios.

Looking at the ATC dataset, when comparing both exploration ratios, we see that the policies behave in a consistent way in both cases. MC-MC and MC-RR present similar results with the 50% ratio, but with 10% ratio the combination that also takes into account the uncertainty to decide when to explore performs slightly better on average after day 25. For the UoL dataset, there is not much improvement in the learning speed, however the MC-MC strategy maintains a slightly lower error than the rest, especially looking at the 50% case.

As a general conclusion, the comparison between MC-MC, MC-RR and RR-MC suggests that the selection of the strategy to decide the location has a higher impact on model accuracy than the temporal policy. While MC-MC and MC-RR obtain similar results, RR-MC results in a decrease in performance compared to MC-MC.

VII. CONCLUSIONS AND FUTURE WORK

This paper proposes an exploration approach for mobile robots which is able to build and refine a spatio-temporal model of pedestrian motion. The exploration strategy uses the uncertainty of an nonhomogenous Poisson process built from past observations to infer the locations and intervals to explore at future times. The experiments show that using the uncertainty measure over the data already gathered, we can increase the learning speed over non-informed strategies,



Fig. 4: Comparison of exploration strategies with partial spatial observations ('where to explore?'). ATC (left) and UoL (right).



Fig. 5: Comparison of exploration strategies with partial temporal observations ('when to explore?'). Dashed and solid lines represent a 10% and a 50% exploration ratio respectively. ATC (left) and UoL (right).



Fig. 6: Comparison of the exploration strategies with both spatial and temporal partial observations ('where and when to explore?'). Dashed and solid lines represent a 10% and a 50% exploration ratio respectively. ATC (left) and UoL (right).

which leads to a better understanding of how the people will behave in certain areas, hence enabling the robot to plans its activities around humans in a more optimal way.

One of the limitations of the method developed lies in the fact that once the number of periodicities m is chosen (in our experiments m = 1), all the cells in the map and the corresponding uncertainties are predicted using the same m. In some scenarios, it is unlikely that all cells in an environment will experience the same rhythmic activities. Future work could therefore analyse the cells individually and choose for each one the best model order to fit the data. Also in this work we are not simulating a realistic behaviour of the robot, omitting many complex actions that would occur in the real world, such as people avoidance, occlusions, battery charging, sensor noise, etc., which could affect the final results and therefore should be investigated further.

ACKNOWLEDGEMENT

This work was supported within H2020-ICT by the EC under grant number 732737 (ILIAD).

REFERENCES

- D. Brscic, T. Kanda, T. Ikeda, and T. T. Miyashita. Person position and body direction tracking in large public spaces using 3d range sensors. *IEEE Transactions on Human-Machine Systems*, 43(6):522–534, 2013.
- [2] Brian Yamauchi. Frontier-based exploration using multiple robots. In Proc. of the Second International Conference on Autonomous Agents, pages 47–53, 1998.
- [3] D. Holz, N. Basilico, F. Amigoni, and S. Behnke. Evaluating the efficiency of frontier-based exploration strategies. In *ISR 2010* (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics), pages 1–8, June 2010.
- [4] Francesco Amigoni and Vincenzo Caglioti. An information-based exploration strategy for environment mapping with mobile robots. *Robotics and Autonomous Systems*, 58(5):684 – 699, 2010.
- [5] J. Vallvé and J. Andrade-Cetto. Dense entropy decrease estimation for mobile robot exploration. In 2014 IEEE International Conference on Robotics and Automation (ICRA), pages 6083–6089, May 2014.
- [6] M. Zreda, W. J. Shuttleworth, X. Zeng, C. Zweck, D. Desilets, T. Franz, and R. Rosolem. Cosmos: the cosmic-ray soil moisture observing system. *Hydrology and Earth System Sciences*, 16(11):4079– 4099, 2012.
- [7] Jaime Pulido Fentanes, Amir Badiee, Tom Duckett, Jonathan Evans, Simon Pearson, and Grzegorz Cielniak. Kriging-based robotic exploration for soil moisture mapping using a cosmic-ray sensor. arXiv preprint arXiv:1811.05384, 2018.
- [8] Cyrill Stachniss and Wolfram Burgard. Mobile robot mapping and localization in non-static environments. In *Conf. on Artificial Intelli*gence, 2005.
- [9] R. Marchant and F. Ramos. Bayesian optimisation for intelligent environmental monitoring. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2242–2249, Oct 2012.
- [10] A. Singh, F. Ramos, H. D. Whyte, and W. J. Kaiser. Modeling and decision making in spatio-temporal processes for environmental surveillance. In 2010 IEEE International Conference on Robotics and Automation, pages 5490–5497, May 2010.
- [11] João Machado Santos, Tomáš Krajník, and Tom Duckett. Spatiotemporal exploration strategies for long-term autonomy of mobile robots. *Robotics and Autonomous Systems*, 88:116 – 126, 2017.
- [12] Tomáiš Vintr, Kerem Eyisoy, and Tomáiš Krajnk. A practical representation of time for the human behaviour modelling. *Forum Statisticum Slovacum*, 14:61–75, 2018.
- [13] R. Senanayake and F. Ramos. Directional grid maps: modeling multimodal angular uncertainty in dynamic environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018. to appear.
- [14] T. Kucner, J. J. Saarinen, M. M. Magnusson, and A. J. Lilienthal. Conditional transition maps learning motion patterns in dynamics environments. In *IEEE International Conference on Intelligent Robots* and Systems, 2013.
- [15] F. Jovan et al. A Poisson-spectral model for modelling temporal patterns in human data observed by a robot. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4013–4018, Oct 2016.
- [16] S.Molina, G.Cielniak, and T. Duckett. Modelling and predicting rhythmic flow patterns in dynamic environments. In *Towards Autonomous Robotic Systems (TAROS)*, volume 10965, pages 135–146, 2018.
- [17] Tomáš Krajník, Jaime Pulido Fentanes, João Santos, and Tom Duckett. FreMEn: Frequency map enhancement for long-term mobile robot autonomy in changing environments. *IEEE Transactions on Robotics*, 2017.
- [18] D. Fink. A compendium of conjugate priors. Technical Report. Cornewell University, 1997.
- [19] Zhi Yan, Tom Duckett, Nicola Bellotto, et al. Online learning for human classification in 3d lidar-based tracking. In *International Conference on Intelligent Robots and Systems (IROS)*, 2017.