

A Poisson-Spectral Model for Modelling Temporal Patterns in Human Data Observed by a Robot

Ferdian Jovan¹, Jeremy Wyatt¹, Nick Hawes¹, and Tomáš Krajník²

Abstract—The efficiency of autonomous robots depends on how well they understand their operating environment. While most of the traditional environment models focus on the spatial representation, long-term mobile robot operation in human populated environments requires that the robots have a basic model of human behaviour.

We present a framework that allows us to retrieve and represent aggregate human behaviour in large, populated environments on extended temporal scales. Our approach, based on time-varying Poisson process models and spectral analysis, efficiently retrieves long-term, re-occurring patterns of human activity from robot-gathered observations and uses these patterns to i) predict human activity level at particular times and places and ii) classify locations based on their periodic patterns of activity.

The application of our framework on real-world data, gathered by a mobile robot operating in an indoor environment for one month, indicates that its predictive capabilities outperform other temporal modelling methods while being computationally more efficient. The experiment also demonstrates that spectral signatures act as features that allow us to classify room types which semantically match with humans’ expectations.

I. INTRODUCTION

Modelling human activities is necessary to succeed in human robot interaction and robot planning of interactions. As a key goal of robots is interaction with human beings, activity models should serve not only to characterize and identify ongoing activities, it should also account for when and where those activities are normally performed. Much robotics research has focused heavily on how to identify activities, leaving predicting “when” and “where” (i.e. spatio-temporal context) of those activities are likely to happen to the experts [1], [2], [3].

Inferring when activities are likely to happen is possible if there are periodic patterns tied to each activity. Fortunately, in human-centered environments, activities exhibit strong rhythmic patterns (daily, weekly, etc). For example, during term time, students normally come to class in the morning, populate canteen areas for lunch, and leave campus before evening. Similar periodic patterns are observable in many types of data such as traffic on a motorway [4], and trading on a stock exchange.

In a similar way to inferring temporal patterns of activities, knowing where activities normally take place is possible if there are definitions of bounded spaces in terms of their functions. These functionalities are normally defined by




Fig. 1: Map of the robot environment and example spectral signatures of the selected rooms. The global map consists of two open plan areas {1,3}, a corridor {12}, a kitchen {5}, three single occupancy offices {2, 10, 11}, three regular offices {4, 8, 9}, and two meeting rooms {6, 7}. Similarity of room 8 and 9 signatures indicates that these rooms have a similar category.

activities being performed in those places, such as offices for working, canteens for having lunch, and kitchens for cooking. Based on their functionality, places which serve the same purpose quite often display the same activities. Hence, they roughly exhibit the same periodic patterns.

In this paper, we focus on the problems of predicting when and where activities are likely to take place, and characterising places according to their activity patterns to maximise human-robot interactions. We simply represent the count of activities performed by humans in each place during some time interval. The temporal pattern of activities is identified by the fluctuating number of humans at particular locations. To learn these temporal patterns in each place, we extend a framework introduced in [5] which is based on Bernoulli distributions.

We replace Bernoulli distributions representing the presence of humans with Poisson processes representing the number of humans in particular places at particular times. Bernoulli distributions, which represent data in binary states, are not suitable for representing the level of activity since they will limit the level of activity to being present or absence. The level of activities provides the information about how busy a place is, hence, leading to how interesting a place is given particular times. Because our aim is to maximise human robot interactions, models with Poisson processes seem to be the better option.

¹Intelligent Robotics Lab, School of Computer Science, University of Birmingham, United Kingdom, fxj345@cs.bham.ac.uk

²Lincoln Centre for Autonomous Systems, University of Lincoln, United Kingdom, tkrajnik@lincoln.ac.uk

This work was funded by the EU ICT project 600623 ‘STRANDS’.

We use the proposed temporal model for two purposes. First, the model can predict the level of human activity in a particular space at a particular time. Second, parameters of the model can act as spatio temporal signatures that allow us to classify the types of individual room. To evaluate the performance of our model in terms of its predictive capabilities, we compare the accuracy of its predictions with a state-of-the-art probabilistic model (Gaussian Process Regression with periodic kernels). To verify the ability to classify location types, we perform hierarchical clustering of the temporal signature and compare the resulting clusters to the real room types.

II. RELATED WORK

There has been recent interest in predicting regular patterns in time series data and using these to find abnormal ones. Several general methods are designed to deal with time series with periodicities, from models such as Autoregressive Moving Average (ARMA) to kernel-based nonparametric models such as Gaussian processes (GPs). Ihler et al. [4], [6] described a modified Markov-modulated Poisson processes for detecting unusual data points or segments in time-series. The Poisson processes are used as probabilistic models for counting regular patterns and behaviour whereas the Markov chain is used to track the occurrence of anomalous events.

Ghassemi and Deisenroth [7] proposed periodic Gaussian processes by re-parametrising the periodic kernel in combination with a double approximation to allow analytic long term forecasting of periodic patterns. Duvenaud et al. [8], [9] introduced a fully automated Bayesian framework based on Gaussian processes with self structured kernel choices, which are built compositionally by adding and multiplying a small number of base kernels. The framework can automatically model any combination of high-level characteristics of time series data, such as smoothness, periodicity and linear trends.

Our framework is derived from the Frequency Map Enhancement (FreMEN) technique proposed by Krajník et al. [5] for spatio-temporal environment representations in long-term scenarios. The FreMEN technique is based on Fourier analysis in combination with a Bernoulli distribution to represent the binary state of data. It has been used in many applications, such as in occupancy grids to compress long-term observations [10], in topological maps to improve robotic search [11], and in path planning [12]. The technique can be applied to all models that represent the world as a set of independent components with binary states [13]. We extend the technique by employing both Poisson processes as the counting model to replace the binary states of FreMEN and a new way of selecting the most prominent frequency components of the Fourier spectrum.

III. DATA SET

Our dataset is a collection of *human trajectories* resulting from a long term deployment of the mobile robot. The data are from a one-month deployment in a building, using a Metralabs Scitos A5 mobile robot equipped with a robust

human tracking algorithm which can detect humans passing within range of its sensors [14]. Since these observations are done by a mobile robot, most of the stored data are incomplete. The detected human trajectories represent only a small fraction of a person's motion.

The partial information collected by a mobile robot is unavoidable because a mobile robot can not fully sense its environment. It can only perceive partial data at a particular time and place. Moreover, the robot's own movement, sensor limitations, and changes in the environment also affect what information a robot can perceive. As a result, our dataset is a collection of chronologically clipped histories about what the robot saw during its observation. Hence, any kind of inference from our dataset is a challenge.

The tracking algorithm we used in our robot produced many false positives, including table legs and chairs. To remove false positives from our dataset, a simple filtering method was used. This is based on the *displacement pose ratio*, which means the distance between the first pose and the last pose of the trajectory over the number of poses in the trajectory. We did not simply remove all short trajectories, having length less than 1m, because information regarding where the persons usually were might be lost. We rather chose to take trajectories with its displacement ratio as the highest ten percent as our dataset. With this filtering, false positives still appear, but the number of them is significantly reduced.

Since the building where our robot was deployed is a large area, we hand-segmented the office into semantic regions such as offices, open plans, a kitchen and corridors. From this process, we obtained 12 datasets, one dataset for each semantic region, over a four-week period. The segmented regions can be seen in the global map in Figure 1.

All collected and filtered human trajectories are used as inputs for the Poisson model. Using Bayesian estimation, we calculate arrival rates for many Poisson processes spread over a month, resulting in a time series of arrival rates. The time series is then analysed via Fourier analysis, to extract its temporal periodic structure. This periodic structure is then used to both predict the level of human activity in a particular space at a particular time and to find types of places forming sensible clusters.

IV. PROBABILISTIC COUNTING MODEL AND SPECTRAL REPRESENTATION

Poisson Models

The appropriate probabilistic model for count data is the Poisson distribution. The probability mass function of the Poisson distribution is:

$$P(N; \lambda) = \frac{e^{-\lambda} \lambda^N}{N!} \quad N = 0, 1, 2, \dots \quad (1)$$

where the parameter λ represents the *rate, or average number of occurrences* in a fixed time interval, and N is the number of occurrences.

Here we refer to $\mathcal{N}(t_i, t_j)$ as a measurement of the number of individuals or objects detected over the time interval




Fig. 2: The λ time series of the corridor updated over 4 week period.

$[t_i, t_j]$ for $i, j \in \{1, \dots, T\}$. We thus transform our λ to be a function of time, i.e. $\lambda(t_i, t_j)$. Hence, (1) becomes a non-homogeneous Poisson distribution, in which the degree of heterogeneity depends on the function $\lambda(t_i, t_j)$. As we use a fixed time interval at any point in time, we define $\lambda(t_i, t_{i+\delta})$ for $i \in \{1, \dots, T\}$ and $\delta \in \mathbb{N}$.

Working with datasets with much missing data means we need a confidence sensitive estimate for the λ . This can be achieved via *Bayesian estimation* to fully calculate the *posterior distribution*. We chose a Gamma distribution conjugate distributions for learning the Poisson parameter λ :

$$\lambda \sim \Gamma(\lambda; \alpha, \beta) \quad (2)$$

The posterior distribution of $\lambda(t_i, t_j)$ given data points x_1, \dots, x_n is calculated as

$$P(\lambda|x_1, \dots, x_n) = \Gamma(\lambda, \alpha + \sum_{i=1}^n x_i, \beta + n) \quad (3)$$

where α, β are the shape and the inverse scale parameter of the Gamma distribution [15].

To fit the Poisson processes and provide the model with good confidence estimate for the λ , we impose one periodicity by splitting the monthly dataset into a weekly period. For each weekly dataset, we calculate the number of trajectories appearing every specified time interval, i.e., every 10 minutes. We then update our Poisson distribution at each time interval. As we use conjugate prior distributions, the rate λ for each time interval is updated by updating the Gamma distribution. The *Maximum a posteriori (MAP)* is chosen to be the point estimate for each updated λ . The point estimate for each λ throughout an entire week creates the λ time series. This is what we refer to as the *Poisson process model*. Figure 2 shows an example of how the λ time series over a week looks after being estimated using a four-week long dataset. The red bar at each point in Figure 2 shows the upper and lower bound of the confidence interval of each λ .

Spectral Representation in Fourier Transform

The *Fourier transform* is a reversible, linear transformation that decomposes a function of time $f(t)$ into the frequencies




Fig. 3: A comparison between the l best amplitude model, the l addition amplitude model, and the Poisson model (the original non periodic model).

that make it up $F(\omega)$. The function $F(\omega)$ is commonly referred to as the frequency spectrum of $f(t)$.

The spectral model - We have shown how we model the occurrence rate λ as a function of time, i.e. $\lambda(t_i, t_j)$. Since we have multiple regions having their own λ time series, we assume that each region is independent to each other. Hence, we can explain the use of the Fourier transform on aa time series of λ for an individual region.

The Fourier transform extracts periodic patterns from λ by calculating the frequency spectrum of λ , i.e. $F(\omega) = FT(\lambda)$. In [5], l coefficients with the highest absolute value along with their frequencies ω_k (for $k = 1, \dots, l$) are selected. For later reference, we call this technique the *l Best Amplitude Model (BAM)*. The coefficients are then used to reconstruct the smoothed signal by means of the inverse Fourier transform $\lambda' = IFT(F'(\omega))$.

Selecting the l best coefficients is a way to filter other frequencies, which are prone to noise, to have a smoother reconstruction signal. However, this technique can not completely capture the magnitude of the original signal whenever the sampling rate is significantly higher than the highest desired frequency. In other words, the higher the ratio between the total number of data points and the highest observed frequency, the smaller the value of the l coefficient with the highest absolute value. Figure 3 shows a signal formed of 30 different periodic signals and stretched over 10000 data points and its reconstruction of l BAM. The highest predefined signal repeats itself 109 times over the data points. It gives the ratio between data points and the highest signal 92.5 against 1. As a result the reconstruction of l BAM technique has a somewhat smaller magnitude than the original signal, even though it captures all the predefined frequencies.

We modified the way to obtain l coefficients in [5] to tackle the aforementioned problem by extracting multi-periodic patterns. To obtain a Fourier spectrum of the raw data, we find a frequency ω_k with the highest absolute value, then subtract it from the data and transform it again. Whenever we obtain a frequency we have encountered, the absolute value is added to the absolute value of the frequency that we have

encountered. We iterate this multiple times until we obtain l desired coefficients. We adopted this technique from [16] applied to get multiperiodic pulsation from observed stars. For later reference, we call this technique the l *Addition Amplitude Model* (AAM). Figure 3 shows how close the AAM reconstruction is to the original Poisson model at each point. This displays that the AAM captures the magnitude of the original Poisson model much better than BAM does.

The results are stored as a set \mathcal{S} consisting of l triples $abs(\omega_k)$, $arg(\omega_k)$, and ω_k which describe the amplitudes, the phase shifts and frequencies of the spectral model. The detailed procedure of l AAM can be seen in Algorithm 1.

Algorithm 1 l addition amplitude model (AAM)

Input: x_1, \dots, x_n : input signal,

total: maximum total frequency

Output: \mathcal{S} : a collection of $(abs(\omega_k), arg(\omega_k), \omega_k)$

Procedure:

1. Init. $k \leftarrow 1$
 - // Get the frequency zero ($\omega_1 = 0$)
 2. $\omega_k \leftarrow FT(x_1, \dots, x_n)[0]$
 3. $\mathcal{S} \leftarrow [[abs(\omega_k), arg(\omega_k), \omega_k]]$
 4. Repeat until $k > \text{total}$
 - $k \leftarrow k + 1$
 - // Get the frequency with the highest amplitude
 - $\omega_k \leftarrow FT(x_1, \dots, x_n)[1]$
 - // Update \mathcal{S} with ω_k
 - if $\omega_k \in \mathcal{S}$, $abs(\mathcal{S}[\omega_k]) \leftarrow abs(\mathcal{S}[\omega_k]) + abs(\omega_k)$
 $arg(\mathcal{S}[\omega_k]) = avg(arg(\omega_k))$
 - else $\mathcal{S} \leftarrow \mathcal{S} + [[abs(\omega_k), arg(\omega_k), \omega_k]]$
 - // Create a cosine signal from ω_k and subtract
 - $x'_1, \dots, x'_n \leftarrow abs(\omega_k) * cos(2\pi * \omega_k + arg(\omega_k))$
 - $x_1, \dots, x_n \leftarrow x_1, \dots, x_n - x'_1, \dots, x'_n$
-

Model purposes - The l addition amplitude model serves two purposes: reconstruction of the original λ time series, i.e., the Poisson processes, and representing spatio-temporal signatures for each region. Reconstructing λ time series is done via the inverse Fourier transform, $\lambda' = IFT(\mathcal{S})$, where \mathcal{S} is a collection of significant periodic patterns of λ extracted via the Fourier analysis.

Examining the Poisson processes and spectral analysis associated with each region, Figure 1 shows that different regions have similar patterns. As each set \mathcal{S} represents the periodic patterns occurring in each region, this enables us to characterize regions according to the similarity of their set \mathcal{S} . Hence, the set \mathcal{S} for each region can be seen as a spatio-temporal signature for that region.

There are several feature sets which can be constructed from the set \mathcal{S} . Here, all frequencies ω_k from each region are put together in the descending order based on how many regions each frequency appears in. m most common frequencies are then selected. These frequencies $\omega_1, \dots, \omega_m$ are then treated as bins where the value of these bins come from $abs(\omega_k)$, or $arg(\omega_k)$, or a combination of both with $k \in \{1, \dots, m\}$.

TABLE I: Comparison of the predictive accuracy of root mean squared error (RMSE) of Poisson model, two Poisson spectral models, and the Automatic Statistician using synthetic datasets.

Method	RMSE	
	no-noise	noise
Poisson processes	101.14 unit	167.98 unit
l Addition amplitude model (AAM)	101.13 unit	161.48 unit
l Best amplitude model (BAM)	109.11 unit	165.53 unit
Automatic statistician (AS - 5 kernels)	101.57 unit	170.31 unit

V. ALGORITHM PERFORMANCE

To investigate the quality of our proposed model in finding periodic patterns from long term observations, we include comparisons with a model with similar capabilities. For this comparison we choose the *Automatic Statistician* (AS) framework of [8], which employs Gaussian processes. Gaussian processes (GP) are a state-of-the-art method for learning models from data. It has been shown that GP with a periodic kernel is able to find repeating patterns for long-term forecasting [7]. Throughout the test, we limited the maximum number of kernel compositions to five. Adding more kernels to the AS model makes the time to construct and calculate the coefficient matrix infeasibly long.

The comparison is based on each model's ability to predict the level of human activity across time and space. Due to the nature of our comparison, we exclude the original FreMEN model from this because FreMEN models the presence of an activity rather than the level of the activity itself [5]. We do compare the original technique for extracting periodic patterns, *BAM* to our current technique *AAM*.

To complete our models for comparison, we include the original non-periodic Poisson processes model. One should note that for the purpose of comparison, we shrink the λ time series to one fifth resulting in fewer data points to fit our dataset for the Automatic Statistician. This is because the implementation could not handle the size of the covariance matrix needed for the Gaussian process [8].

At the end of the section, we study the ability of the best method to classify rooms or regions according to their spatio-temporal signatures. Two different clustering algorithms are presented to show as a comparison.

Validation on Synthetic Data

First we validated the ability of different models to recover periodic patterns on a set of synthetic data. The synthetic dataset was created from 30 different periodic patterns. We then added Gaussian noise to each point in the synthetic data.

Our synthetic dataset follows the format of our four-week long real dataset. We performed four-fold-cross-validation (CV) on the synthetic dataset where each CV-fold is a different week. We compared four models including the Poisson processes, AAM, BAM, and AS. We record the root mean squared error (RMSE) of the reconstructions in Table I.

With the absence of noise, Table I shows that many models performed similarly, with the exception of BAM which

TABLE II: Comparison of the predictive accuracy of root mean squared error (RMSE) of Poisson model, the Poisson spectral models, and the Automatic Statistician using real-world dataset.

Method	RMSE for each region												Average
	1	2	3	4	5	6	7	8	9	10	11	12	
Poisson model	2.63	9.50	7.32	2.65	5.85	2.92	4.40	4.13	1.91	1.54	4.93	5.02	4.40
<i>l</i> AAM	2.49	8.99	6.61	2.65	5.66	2.45	4.20	4.04	1.87	1.40	4.83	4.67	4.15
<i>l</i> BAM	2.57	9.02	6.93	2.65	5.69	2.45	4.20	4.04	1.87	1.40	4.82	4.95	4.22
AS (5 kernels)	2.49	8.67	6.71	2.66	5.76	3.40	4.34	4.02	1.97	1.47	5.03	4.79	4.27

TABLE III: Comparison of the learning time of the Poisson spectral models and the Automatic Statistician using real-world datasets. Note that the automated statistitian times are in hours, while the AAM and BAM times are in seconds.

Method	Learning time for each region												Average
	1	2	3	4	5	6	7	8	9	10	11	12	
<i>l</i> AAM	1.1s	1.0s	1.3s	0.8s	0.8s	0.6s	0.8s	1.0s	0.8s	1.0s	1.2s	1.4s	0.9s
<i>l</i> BAM	0.1s	0.2s	0.1s	0.1s	0.2s	0.1s	0.1s	0.2s	0.1s	0.3s	0.2s	0.1s	0.2s
AS (5 kernels)	3.4h	1.8h	4.0h	1.9h	2.5h	2.4h	1.4h	2.0h	2.0h	1.2h	2.9h	1.5h	2.3h

performed poorly. In the presence of noise, AAM and BAM showed their competence by outperforming any other model. The table clearly shows that models with spectral analysis performed slightly better than a fully Bayesian approach in the presence of noise.

One should note that we used strong uniform priors for our Poisson processes which are suitable for our real world datasets. Our priors are based on the assumption that people appear in any time of a day is unlikely to happen. In other words, the arrival rate λ at any time interval is close to zero. We did not try to find suitable priors to match our synthetic dataset. As a result, the Poisson processes did not perform really well on our synthetic dataset with the average error 100 units without noise and 167 units with noise. Nonetheless, this does not affect the relative performance of our reconstruction model which is slightly better than AS reconstruction since both of the reconstructions are based on Poisson processes.

Performance on Real World Datasets

We compared the four models described in the previous section in terms of their predictive accuracy on our real-world datasets. We performed four fold cross-validation in a weekly manner on the collected datasets as described in Section III.

Results are presented in Table II. From the average result, AAM, BAM, and the AS model improved the predictive accuracy of the Poisson processes by 6%, 4.3%, and 3.1% respectively. Once more, the table confirms that in the presence of noise, models with spectral analysis performed slightly better than Gaussian Processes. Moreover, Table I and II show that the introduction of the AAM technique for periodic pattern extraction improves the predictive accuracy of our reconstruction model.

We also present the time needed for the learning reconstruction (Table III). Our finding here is based on our reduced dataset, explained earlier in this section. In terms of speed, BAM outran other models at least by a factor of 5. The AAM

is still fast (1 second on average). This leaves the AS model by far the slowest one, with at least one hour needed to learn.

A. Clustering Capability

To test the hypothesis that different regions have similar patterns, we consider a clustering approach. We require a clustering process that makes weak prior assumptions about the number of room classes and which will produce a hierarchical structure capturing the room similarities. For this we employ Dirichlet Process (DP)-means clustering [17]. This algorithm combines Dirichlet process mixture models and a classical clustering algorithm, so as to have scalable algorithms that retain the main benefit of Bayesian non-parametrics, which is the ability to model infinite mixtures. Using this clustering algorithm, we range over the penalty parameter rather than explicitly deciding the number of clusters prior to the learning process. We compared this to the standard K-Means algorithm.

Using the AAM model, each clustering process constructs a tree which expresses the similarities between room types in a hierarchical fashion. For the DP-means clustering, the dendrogram was produced by varying the penalty parameter, whereas for K-Means, the dendrogram was produced by varying the number of clusters we would like to have. Figure 4 shows the dendrograms produced by DP-Means clustering (4a) and K-Means clustering (4b).

From Figure 4, it is easy to verify that clusters produced by two clustering algorithms are sensible. Those clusters can be used to represent the general function/type of a room. Moreover, the clustering hierarchy of the algorithms matches with a semantic room type hierarchy. One should note that the single occupancy office $\{2\}$ is a special case. This room belongs to the manager. It thus has quite a different pattern of activity.

VI. CONCLUSIONS

We have presented an approach to building a probabilistic model of time-varying counting processes. We have shown




Fig. 4: Dendrograms of region clustering.

that this can find regular (periodic) patterns in human behaviour. The approach is based on an assumption that aggregate statistics of human activities have periodicities which can be observed from the fluctuating number of humans around. These periodic patterns can be described by means of frequency, amplitude, and phase, modelled using the Fourier and inverse Fourier transforms. By taking the most significant spectrum components of the Fourier transform, we indirectly obtain the most significant periodic patterns in the human activity level. As each region might have a unique frequency spectrum, the spectrum components can be further used as features for region-type clustering.

We then evaluated the performance of the proposed framework on several time-series of counts representing tracked people, which were collected by an autonomous mobile robot in an indoor environment over a month. The results indicate that the proposed framework is able to produce the model up to 1000 times faster than the Automatic Statistician framework with competitive prediction performance. Moreover, we demonstrated that the spectral representation of the model serves a dual purpose by allowing us to cluster regions by their spatio-temporal signatures. The clusters produced by our framework show an intuitive result in which the clusters match roughly with human expectations of room-type clustering.

In this paper, we have performed the temporal analysis independently for each room. An interesting extension would be to automatically understand the relationship between the time series for all different rooms. Furthermore, we will investigate the descriptiveness of other methods with our spatio-temporal signatures in room-type clustering.

REFERENCES

- [1] M. Sridhar, A. G. Cohn, and D. C. Hogg, "Unsupervised learning of event classes from video," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. AAAI Press, 2010, pp. 1631–1638.
- [2] —, "Discovering an event taxonomy from video using qualitative spatio-temporal graphs," in *ECAI 2010-19th European Conference on Artificial Intelligence, Proceedings*, vol. 215. IOS Press, 2010, pp. 1103–1104.
- [3] K. S. R. Dubba, A. G. Cohn, and D. C. Hogg, "Event model learning from complex videos using ilp," in *ECAI*, vol. 215, 2010, pp. 93–98.
- [4] A. Ihler, J. Hutchins, and P. Smyth, "Adaptive event detection with time-varying poisson processes," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 207–216.
- [5] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, "Spectral analysis for long-term robotic mapping," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 3706–3711.
- [6] A. Ihler, J. Hutchins, and P. Smyth, "Learning to detect events with markov-modulated poisson processes," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 3, p. 13, 2007.
- [7] N. H. Ghassemi and M. P. Deisenroth, "Analytic long-term forecasting with periodic gaussian processes," in *Proc. of AISTATS*, 2014, pp. 303–311.
- [8] D. Duvenaud, J. R. Lloyd, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani, "Structure discovery in nonparametric regression through compositional kernel search," in *Proceedings of the 30th International Conference on Machine Learning*, June 2013, pp. 1166–1174.
- [9] J. R. Lloyd, D. Duvenaud, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani, "Automatic construction and Natural-Language description of nonparametric regression models," in *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [10] T. Krajník, J. Santos, B. Seemann, and T. Duckett, "Froctomap: an efficient spatio-temporal environment representation," *Advances in Autonomous Robotics Systems*, p. 269, 2014.
- [11] T. Krajník, M. Kulich, L. Mudrova, R. Ambrus, and T. Duckett, "Where's waldo at time t? using spatio-temporal models for mobile robot search," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 2140–2146.
- [12] J. P. Fentanes, B. Lacerda, T. Krajník, N. Hawes, and M. Hanheide, "Now or later? predicting and maximising success of navigation actions from long-term experience," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1112–1117.
- [13] T. Krajník, J. Pulido Fentanes, J. Santos, K. Kusumam, T. Duckett, *et al.*, "Fremen: frequency map enhancement for long-term mobile robot autonomy in changing environments," in *ICRA15 WS VPRiCE*. IEEE, 2015.
- [14] C. Dondrup, N. Bellotto, F. Jovan, M. Hanheide, *et al.*, "Real-time multisensor people tracking for human-robot spatial interaction," in *Workshop on Machine Learning for Social Robotics at International Conference on Robotics and Automation (ICRA)*. IEEE, 2015.
- [15] D. Fink, "A compendium of conjugate priors," *Technical report, Cornell University*, 1995.
- [16] M. Templeton, "Time-series analysis of variable star data," *Journal of the American Association of Variable Star Observers (JAAVSO)*, vol. 32, no. 1, pp. 41–54, 2004.
- [17] B. Kulis and M. I. Jordan, "Revisiting k-means: New algorithms via bayesian nonparametrics," in *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012, pp. 513–520.