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

Ferdian Jovan[1], Jeremy Wyatt[1], Nick Hawes[1], and Tomáš Krajník[2]

***Abstract*— The efficiency of autonomous robots depends on how well they understand their operational 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 understanding 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 combination of 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 predict human activity and classify locations.

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 traditional 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 the "when" and "where" (i.e. spatio-temporal context) of those activities to the experts. In this paper, we instead focus on the problem of predicting when and where an activity is likely to take place, and to characterise places according to their activity patterns. This knowledge can be used to drive robot interaction with humans.

One insight is that human activities rhythmic spatio-temporal patterns. In particular, some of the activities are periodic on a number of scales (daily, weekly, etc). For example, the rhythm of making coffee is different from one person to another. Some might want a cup of coffee once a day, others may want more. Some might want it in the morning, others might want it after lunch. In particular rhythms of aggregated behaviours of a population have strong periodicities. These kinds of rhythms create aggregated behaviours of the population. These rhythms can be detected by the fluctuating number of humans at the location where the particular activity is performed. An example is the level of human activities in an office: it is high during office hours and decreases at the end of the day. The pattern repeats for five consecutive days, and stops during weekends creating a repeating daily/weekly routine. Similar periodic patterns are observable in many types of data such as traffic on a motorway [1], activities in a school, and trading on a stock exchange.

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

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

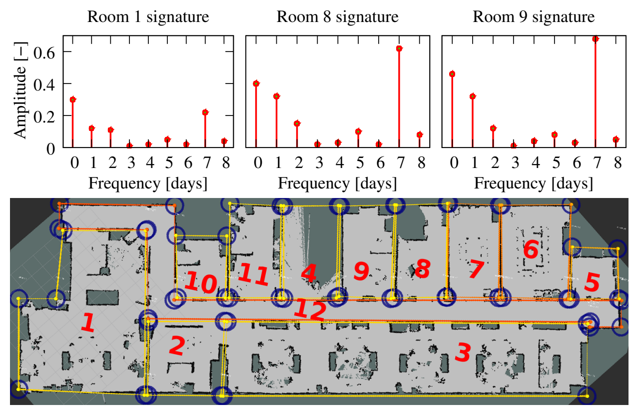

Fig. 1: Map of the robot environment and example spectral signatures of the selected rooms. The global map consists of two open plan areas {1,3}, a corridor {12}, a kitchen {5}, three single occupancy offices {2, 10, 11}, three regular offices {4, 8, 9}, and two meeting rooms {6, 7}. Similarity of room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

In this paper, we present a method that learns the periodic pattern of aggregated human activities within a space by means of Poisson processes combined with a frequency analysis. We focus on time-series count data where time is discrete and $\mathcal{N}(t_i, t_j)$ is a measurement of the number of individuals or objects detected over the time interval $[t_i, t_j)$, e.g counts of the number of people who enter a building every 10 minutes. As this type of data measures the aggregate behaviour of many individuals, it typically exhibits temporal periodicities. We show how to learn such models from data to both characterise regular periodic patterns and detect irregular ones within a time interval. We adopt a Bayesian approach to learning those models.

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 locations. To evaluate the performance of our model in terms of their predictive capabilities, we compare the accuracy of its prediction with the state-of-the-art principled probabilistic model. 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 and finding irregular ones in time series data. 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. [1], [2] 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 [3] 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. [4], [5] 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 characteristic of time series data such as smoothness, periodicity and linear trends.

Our framework is derived from the Frequency Map Enhancement (FreMEn) technique proposed by Krajnik et al. [6] 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 [7], in topological maps to improve robotic search [8], and in path planning [9]. The technique can be applied to all models that represent the world as a set of independent components with binary states [10]. 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 [11]. Similar to other datasets with long term observations, it has a large amount of stored data. However, as these observations are made by a mobile robot, most of the stored data are incomplete. Many of 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 there is a limitation on how much information a robot can perceive. A mobile robot is not an omnipresent being; it can not sense the full 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. In the attempt to remove false positives from our dataset, a simple filtering method was used. This filtering 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 the best ten percent of trajectories, based on the displacement pose ratio, 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 segmented the office into semantic regions such as offices, open plans, a kitchen and corridors. The segmentation represents the actual imaginary segments of the office. 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 will be used as inputs for the Poisson model. Using Bayesian estimation, we calculate arrival rates for Poisson processes over a month period 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 frequency of human activity in a particular space at a particular time and to classify 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:

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

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 $[t_i, t_j)$ for $i, j \in \{1, \ldots, T\}$. We thus transform our λ to

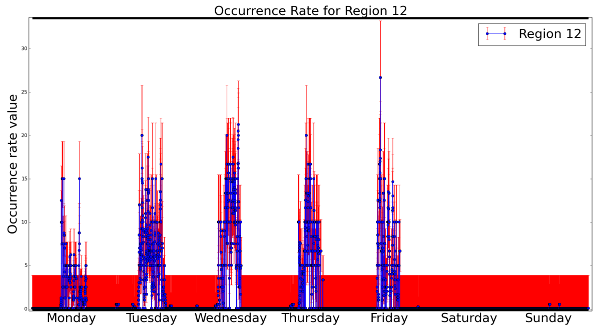

Fig. 2: The λ time series of the corridor updated over 4 week period.

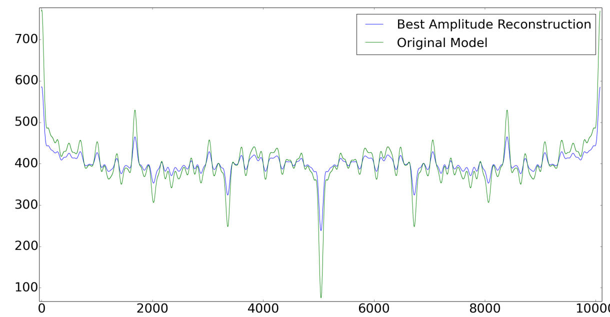

Fig. 3: A comparison between the l best amplitude model and the original data points.

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 $i \in \{1, \ldots, T\}$ and $\delta \in \mathbb{N}$.

Working with datasets with much missing data means we need confidence estimate for the λ. This can be achieved via *Bayesian estimation* to fully calculate the *posterior distribution*. We chose conjugate prior distributions for learning the λ variable to ensure that the inference computations have a simple closed form:

$$\lambda \sim \Gamma(\lambda; \alpha, \beta) \tag{2}$$

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

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

where α, β are the shape and the inverse scale parameter of the Gamma distribution [12].

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 by the *Poisson processes model*. Figure 2 shows an example of how the λ time series over a week looks after being updated by four-week dataset.

Spectral Representation in Fourier Transform

The *Fourier transform* is a reversible, linear transformation that decomposes a function of time $f(t)$ into the frequencies 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 Fourier transform on λ time series of each region.

The Fourier transform is proposed to mine periodic patterns of λ by calculating the frequency spectrum of λ, i.e. $F(\omega) = FT(\lambda)$. In [6], l coefficients with the highest absolute value along with their frequencies ω_k are selected. For later reference, we call this technique *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 frequency that you would like to observe. In other words, the higher the ratio between total 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 somewhat smaller magnitude than the original signal even though it captures all the predefined frequencies.

We modified the way to obtain l coefficients in [6] to tackle the aforementioned problem by mining 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 [13] applied to get multiperiodic pulsation from observed stars. For later reference, we call this technique *l addition amplitude model (AAM)*. The results are stored as a set of $\mathcal{S}$ 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: $x_1, \ldots, x_n$: input signal,
total: maximum total frequency
Output: $\mathcal{S}$: a collection of $(abs(\omega_k), arg(\omega_k), \omega_k)$
Procedure:
1. Init. k = 1
// Get the frequency zero ($\omega_1 = 0$)
2. $\omega_k = FT(x_1, \ldots, x_n)[0]$
3. $\mathcal{S} = [[abs(\omega_k), arg(\omega_k), \omega_k]]$
4. Repeat until k > total
• $k = k + 1$
// Get the frequency with the highest amplitude
• $\omega_k = FT(x_1, \ldots, x_n)[1]$
// Update $\mathcal{S}$ with ω_k
• if $\omega_k \in \mathcal{S}$, then $abs(\omega_k) += abs(\omega_k)$ and $arg(\omega_k) = avg(arg(\omega_k))$
else $\mathcal{S} = \mathcal{S} + [[abs(\omega_k), arg(\omega_k), \omega_k]]$
// Create a cosine signal from ω_k
• $x'_1, \ldots, x'_n = abs(\omega_k) * cos(2\pi * \omega_k + arg(\omega_k))$
// Subtract current $x_1, \ldots, x_n$ with the cosine signal
• $x_1, \ldots, x_n = x_1, \ldots, x_n - x'_1, \ldots, x'_n$

Model purposes - The l addition amplitude model serves for 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 inverse Fourier transform, $\lambda' = IFT(\mathcal{S})$, resulting λ' whose the magnitude is smaller than the original one. Having a smaller magnitude in λ' acts to further reduce noise which is unfiltered during the data preprocessing producing cleaner rates with smooth transitions.

Examining the Poisson processes associated with each region, it can be seen 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 the number of their appearance in all regions. This means the frequency which appears in most of the regions will likely be put in the very beginning. m most common frequencies are then selected. These frequencies $\omega_1, \ldots, \omega_m$ 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 $k \in \{1, \ldots, m\}$.

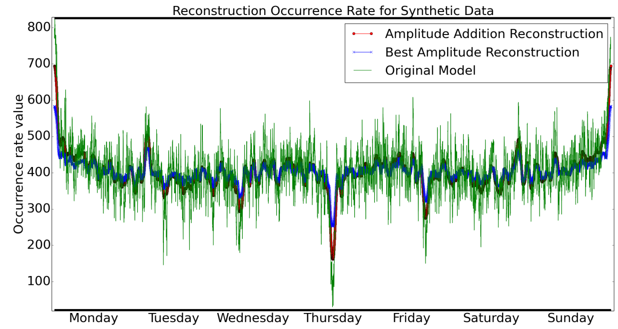

Fig. 4: The comparison between synthetic data and the two reconstruction model.

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

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

V. Algorithm Performance

To know the quality of our model, we analyse the performance and compare it with the automatic statistician framework of [4], which employs Gaussian processes as well as with the original BAM model of [6]. For later reference, we call the automatic statistician framework *automatic statistician (AS)*. We compare the algorithms on their ability to predict the level of human activity across time and space. Following this, we study the ability of the best method to classify rooms or regions according to their spatio-temporal signature. 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 i.i.d Gaussian noise to each point in the synthetic data. Figure 4 shows the Poisson processes of our synthetic dataset with two reconstruction models, the BAM reconstruction following the technique presented in [6] and our AAM reconstruction. The original model shown in the figure is purely the Poisson processes of the synthetic dataset. Figure 4 shows that at each point AAM has smaller distance with the Poisson processes than what BAM has. This shows that AAM captures the magnitude of the Poisson processes better than BAM does.

As the synthetic dataset follows the format of our real dataset, we performed 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,

TABLE II: Comparison of the predictive accuracy according to a root mean squared error (RMSE) measure Poisson model. We compare the Poisson spectral model, and 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
l 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**
l 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 model and automatic statistician using real-world datasets. Note that the automated statistitian times are in hours, while the AAM and BAM in seconds.

Method	Learning time for each region												Average
	1	2	3	4	5	6	7	8	9	10	11	12	
l 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
l 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

and AS. We record the root mean squared error (RMSE) of the reconstructions in Table I.

Table I clearly shows that AAM outperformed others. With the absence of noise, AAM is 7.9% more accurate than BAM. On the other hand, the AS model produced a competitive prediction even though 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. In the presence of noise, AAM and BAM outperformed the Poisson processes. They also outperformed the AS which performed the worst in the presence of noise. Furthermore, Table I shows that AAM is 5.5% more accurate than the AS.

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 with our synthetic dataset with the average error around 100 points for synthetic dataset without noise, and 167 points for synthetic dataset with noise. Nonetheless, this does not affect the relative performance of our reconstruction model which is slightly better than AS reconstruction model since both of the reconstructions are based on the Poisson processes.

Performance on Real World Datasets

We also compared the four models described in the previous section in terms of their predictive accuracy. 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. In general, the AAM outperformed AS in almost every test except for Region 8, which is a regular office, and Region 2, which is a single occupancy office. Table II also shows the average predictive performance among regions. 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. One should note that for the purpose of comparison, we shrink the λ time series to one fifth resulting fewer number of data points to fit our dataset to the automatic statistician. This is because of computational limitations which could not handle the size of the covariance matrix needed for Gaussian learning process.

We also present the time needed for each learner (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 construct a model.

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[14]. This algorithm combines Dirichlet process mixture models and classical clustering algorithm to have scalable algorithms that retain the main benefit of Bayesian nonparametrics, 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 AAM model, each clustering process constructs a tree which expresses the similarities between room types on a hierarchical fashion. For the DP-means clustering, the dendogram was produced by varying the penalty parameter, whereas for the K-Means, the dendogram was produced by varying the number of clusters we would like to have. Figure 5 shows the dendograms produced by DP-Means clustering

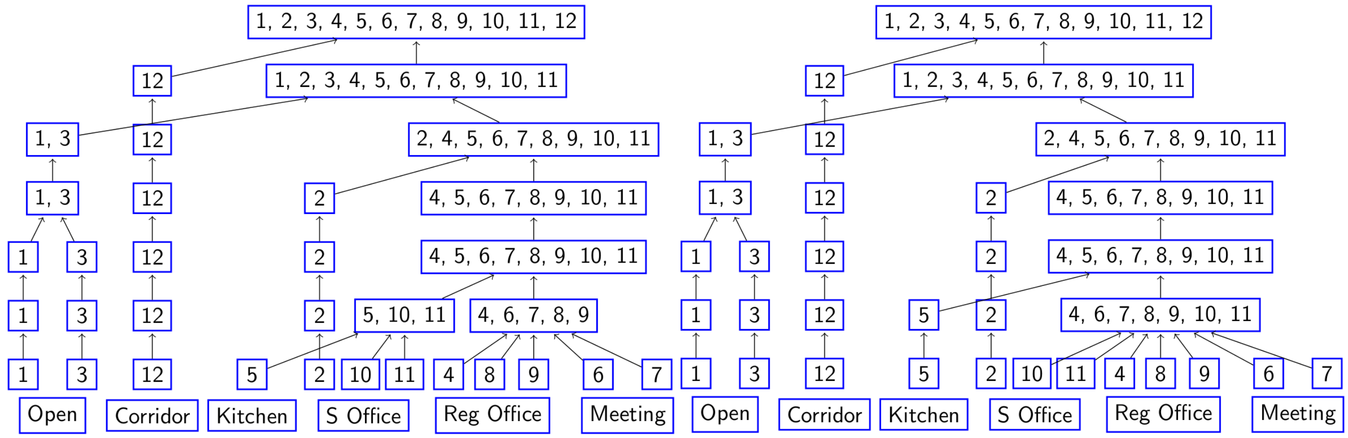

(a) DP-Means clustering.

(b) K-Means clustering.

Fig. 5: Dendograms of region clustering.

(5a) and K-Means clustering (5b).

From Figure 5, 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 rooms. Moreover, the clustering hierarchy of the algorithms matches with semantic room type hierarchy that humans expect. One should note that a single occupancy office $\{2\}$ is a special case of those offices. This room belongs to a person with the highest position in the office there.

VI. Conclusions

We have presented an approach for building a probabilistic model of time-varying counting processes. We have shown 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 that are influenced by the underlying human activities. 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 a competitive predictive measure. 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 at which the clusters match with humans' expectation 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.

References

[1] A. Ihler, J. Hutchins, and P. Smyth, "Adaptive Event Detection with Time-Varying Poisson Processes," *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '06*, p. 207, 2006. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1150402.1150428

[2] ——, "Learning to detect events with markov-modulated poisson processes," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 3, p. 13, 2007.

[3] N. H. Ghassemi and M. P. Deisenroth, "Analytic long-term forecasting with periodic gaussian processes," in *Proc. of AISTATS*, 2014, pp. 303–311.

[4] 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.

[5] J. R. Lloyd, D. Duvenaud, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani, "Automatic construction and Natural-Language description of nonparametric regression models," in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2014.

[6] T. Krajnik, 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.

[7] T. Krajnık, J. Santos, B. Seemann, and T. Duckett, "Froctomap: an efficient spatio-temporal environment representation," *Advances in Autonomous Robotics Systems*, p. 269, 2014.

[8] T. Krajnik, 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.

[9] J. Pulido Fentanes, B. Lacerda, T. Krajnik, N. Hawes, M. Hanheide, *et al.*, "Now or later? predicting and maximising success of navigation actions from long-term experience," 2015.

[10] T. Krajnik, J. Pulido Fentanes, J. Santos, K. Kusumam, T. Duckett, *et al.*, "Fremen: frequency map enhancement for long-term mobile robot autonomy in changing environments," 2015.

[11] C. Dondrup, N. Bellotto, F. Jovan, M. Hanheide, *et al.*, "Real-time multisensor people tracking for human-robot spatial interaction," 2015.

[12] D. Fink, "A compendium of conjugate priors," *See http://www. people. cornell. edu/pages/df36/CONJINTRnew% 20TEX. pdf*, p. 46, 1997.

[13] 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.

[14] B. Kulis and M. I. Jordan, "Revisiting k-means: New algorithms via bayesian nonparametrics," *arXiv preprint arXiv:1111.0352*, 2011.