



Efficiently exploring for human robot interaction: partially observable Poisson processes

Ferdian Jovan¹ · Milan Tomy² · Nick Hawes³ · Jeremy Wyatt⁴

Received: 20 January 2022 / Accepted: 5 October 2022 / Published online: 28 October 2022
© The Author(s) 2022

Abstract

Consider a mobile robot exploring an office building with the aim of observing as much human activity as possible over several days. It must learn where and when people are to be found, count the observed activities, and revisit popular places at the right time. In this paper we present a series of Bayesian estimators for the levels of human activity that improve on simple counting. We then show how these estimators can be used to drive efficient exploration for human activities. The estimators arise from modelling the human activity counts as a partially observable Poisson process (POPP). This paper presents novel extensions to POPP for the following cases: (i) the robot's sensors are correlated, (ii) the robot's sensor model, itself built from data, is also unreliable, (iii) both are combined. It also combines the resulting Bayesian estimators with a simple, but effective solution to the exploration-exploitation trade-off faced by the robot in a real deployment. A series of 15 day robot deployments show how our approach boosts the number of human activities observed by 70% relative to a baseline and produces more accurate estimates of the level of human activity in each place and time.

Keywords Bayesian estimators · Poisson processes · Exploration-exploitation · Human robot interaction

1 Introduction

Autonomous mobile robots are being developed to operate in human populated spaces, such as homes and offices (Hawes et al., 2016). The biggest benefits of mobile robots operating in these spaces over any other intelligent system such as smart houses and assisted living systems are the mobility and the ease of deployment since there is no need to modify existing environment structures. Since, the robot is in human-centered environments, it is useful for these robots to predict patterns of human activity, so as to learn about those activities (Coppola et al., 2016; Hogg & Cohn, 2016) or to plan interactions with humans (Street et al., 2020). This paper is concerned with how a robot can correctly estimate how many human activities it has encountered, predict how many will

occur in a particular location at a particular time, and then use these to optimize the exploration-exploitation trade-off that occurs during active learning, so as to observe as many human activities as possible during a deployment.

Consider a mobile service robot that works in a large office building. Let us suppose that one of the system designers' aims is for the robot to observe and thereby learn about the various activities performed by humans, and for this learning to have to occur over several days or weeks. To achieve this goal, the robot must learn where and when people are to be found, count the observed activities (perhaps grouped by their categories), and revisit places at times when those activities can be observed in sufficient number. For example, to learn about eating activities the robot would benefit from visiting the canteen at lunchtime, rather than at the start of the day. If such a robot is to be deployed to a variety of buildings, without re-programming of rules by hand, it should autonomously learn the spatio-temporal distribution of these activities and exploit that learning to observe a useful variety and number of activities.

This involves solving two problems. First, the robot must estimate where it can observe the greatest number of human activities at each time of the day. Second, as the robot is learning it should trade-off exploring for new time-place

✉ Ferdian Jovan
ferdian.jovan@bristol.ac.uk

¹ Digital Health, University of Bristol, Bristol, United Kingdom

² Rio Tinto Centre for Mine Automation, University of Sydney, Sydney, Australia

³ Oxford Robotics Institute, University of Oxford, Oxford, United Kingdom

⁴ Robotics AI, Amazon, Boston, United States




Fig. 1 Our mobile robot observes people at an event

combinations where it might discover a high-level of human activity and re-visiting those time place combinations where it already knows that a wealth of human activity is to be found. This second problem involves solving an exploration-exploitation trade-off.

This paper presents a Bayesian method to solve both these problems in the case where we treat human activities as count data. The Bayesian framework models not only the frequency of human activities and the variation in this, but also the robot's uncertainty about the mean rate at which activities occur. Thus, the Bayesian estimator captures both inherent process uncertainty (aleatoric uncertainty) and the robot's additional uncertainty in what it knows about the process (epistemic uncertainty). It can also correct for inherent biases (a tendency to false positives or false negatives in the sensory system). Because of this it has two advantages over a baseline frequentist estimator. First, it will produce more accurate estimates and predictions of human activity levels than a method that does not model classification errors. Second, because it captures epistemic uncertainty it can be used to perform active learning. This active learning problem is fundamentally an exploration-exploitation trade-off. Should the robot visit a place at a time such that it can exploit what it already knows about the likely activity level, or should it explore a place-time combination about which it knows less, but about which it might learn and so lead to a higher rate of activity observation in the long run? This active learning problem is intractable in the strict formulation, since it involves reasoning over a tree of possible knowledge states. Despite this, there are effective, heuristic active-learning rules that are quick to evaluate (Fig. 1).

We develop a series of Bayesian estimators. Then we present a method to use these to drive exploration. This uses both a Fourier transform to capture the periodicity of human activities and the epistemic uncertainty in the activity rate as captured by the posterior. Using this estimation, prediction and exploration technique, we then present the results of sev-

eral long-run deployments of a real robot in a public building. These long-run deployments (15 days per treatment) are used to test whether the different Bayesian estimators, together with the solution to the exploration-exploitation trade-off, result in the robot observing greater numbers of human activities than a baseline frequentist method.

This paper builds on our earlier work, which showed how to count reliably from a single unreliable detector or from multiple, unreliable, uncorrelated detectors (Jovan et al., 2018). That work formulated the problem as Bayesian inference for a *partially observable Poisson process* (POPP) and showed an improvement on a baseline model assuming sensor reliability, termed the fully observable Poisson process (FOPP).

This paper makes the following technical contributions. First, we extend the POPP model to create the correlated POPP (C-POPP) model. This supports inference when the robot has multiple detectors with correlated outputs. Second, the observation model used to correct counts in the POPP model is itself constructed from data and so has both epistemic and aleatoric uncertainties. The POPP and C-POPP models only take account of the aleatoric uncertainty in the observation model. We extend the POPP model to include the epistemic uncertainty, resulting in the POPP-Beta model. The third contribution is to combine the benefits of C-POPP and POPP-Beta. This results in the POPP-Dirichlet model, which works for correlated sensors and epistemic uncertainty in the observation model. We demonstrate the inferential properties of POPP and these three extensions in both numerical simulations. The fourth contribution is that we show how these models can be used solve the exploration-exploitation problem by combining Fourier transform that allows us to exploit the periodicity of human activities with an upper bound estimate derived from the posterior. Finally, the fifth contribution is an extensive real world evaluation on a long-run robot. We compare the exploration and estimation performance of the FOPP, POPP and POPP-Beta models in a series of three 15-day deployments. Analysis shows that the POPP and POPP-Beta models are able to explore more efficiently, encountering more people than the baseline FOPP model and that they produce superior estimates of the rate of human activities.

2 Related work

As requirements to employ the FOPP model are unlikely to be met, some existing works on statistical models propose a way to work with observation data that are not fully observable. In some literature, this is termed *misclassified counts*. Misclassification happens when there are *false positive counts* or *false negative counts* (or both). False positive counts, also called the overcount, occur when the count includes events

other than those of interest. False negative counts, also called the undercount, occur when some of the events of interest are missed. Work on the undercounting problem is common. Whittemore and Gong estimated cervical cancer rates by taking into account false negative data (Whittemore & Gong, 1991). Winkelmann and Zimmermann introduced a combination of a Poisson regression model with a logit model for under-counting, yielding the Poisson-Logistic (Pogit) model (Winkelmann & Zimmermann, 1993). They applied this to model the number of days employees were absent from a workplace. Dvorzak and Wagner adapted the Pogit model to use a small set of validation data, to provide information about the true counts (Dvorzak & Wagner, 2016). They performed a Bayesian analysis of the Poisson-Logistic model and incorporate Bayesian variable selection to identify regressors with a non-zero effect and also to restrict parameters of the Poisson-Logistic model.

There is less prior work on the Poisson model for the case where the data may either be undercounted or overcounted (Sposto et al., 1992; Bratcher & Stamey, 2002; Stamey et al., 2004; Stamey & Young, 2005). Sposto et al. followed a frequentist approach to estimate both cancer and non-cancer death rates, assuming false negatives are possible on both sides of these counts (Sposto et al., 1992). In (Bratcher & Stamey, 2002), Bratcher and Stamey used a Bayesian method to estimate Poisson rates in the presence of both undercounts and overcounts, borrowing the double sampling technique introduced in (Tenenbein, 1970). They extended their work to a fully Bayesian method for interval

prediction of the unobservable actual count in future samples, given a current double sample (Stamey et al., 2004). Stamey and Young (Stamey & Young, 2005) present closed-form expressions for maximum likelihood estimators of the false negative rate, the false positive rate, and the Poisson rate for the model proposed in (Bratcher & Stamey, 2002). The estimators are straightforward to calculate and to interpret in terms of evaluating the effectiveness of using unreliable counts.

What we propose is similar to that of Bratcher and Stamey (Bratcher & Stamey, 2002). Both aim to accurately estimate the arrival rate parameter of a single Poisson process. Bratcher and Stamey utilise double sampling to obtain the true count together with false positive and false negative counts. They estimate the rate via MCMC since no closed form is found for λ , and the calculation of the full posterior is expensive. Double sampling assumes access to two counters with one always being a perfect counter. Our work goes beyond this since we consider multiple, potentially correlated, but always unreliable counters. We extend the work of Jovan et al in (Jovan et al., 2018) by presenting three extensions of our original model.

In Sect. 8 we validate our work by demonstrating how it can be used to improve mobile robot exploration missions to observe humans. Existing work in this field is typically driven by the entropy in the model of a process to maximise the outcome of interests from an exploration (e.g. the size of area explored, the number of observed humans). However, many existing works rely on the assumption that the




Fig. 2 A cycle process from count data collection, Poisson process estimation, through exploration plan generation for one day in an office-like environment. Count data are collected through perception algorithms or sensors while the robot patrols. The raw count data from multiple sensors are correctly filtered and merged via Bayesian inference taking

into account the unreliability and correlation of each sensor to estimate the underlying Poisson process on each region of interest. Each estimate of the Poisson process is used by the exploration planner to find the maximum upper-bound of the Poisson process at each time interval and the corresponding region is chosen as a place to visit

sensors are fully reliable and data is fully observable; the collected data is, relatively, free from inherent biases. Molina et al. (2022) is a prime example of this work with the state of the art result in robot explorations. Molina et al. (2022) focuses on robot explorations in learning human motion patterns by (1) exploiting the temporal aspects of human motion through spectral analysis to decide when to visit particular cells and (2) incorporating entropy calculation in the probabilistic maps to decide which places/cells worth visiting. Kaplow et al. (2010) employed a variable resolution map in combination with POMDP formulation to achieve scaling with a robotic wheelchair navigation and exploration. Task-level robot control with a decision-theoretic framework was first tackled by Pineau et al. in (Pineau et al., 2003) using a POMDP planner to derive a high-level controller for a mobile robot with a dialogue system by exploiting hierarchy to reduce the state space.

There are couple of works that do reason about sensing reliability in an attempt to correct biases in the sensors with the application on robot explorations. Martinez-Cantin et al. (2009) give a POMDP formulation of active visual mapping, use direct policy search to find a solution, and use Monte Carlo simulation to generate imaginary observations and action outcomes during optimization. The main challenge of decision-theoretic planning in partially observable environments is intractability. Velez et al. (2012) planned trajectories in a continuous space to maximize the reliability of object detection using a learned observation model. The key contribution is the use of a model of the correlations in sensor behaviour at nearby locations, thus driving the robot to gather more informative views. Similar to the work of Velez et al (Velez et al., 2012) in utilizing sensor behaviours in driving the robot to gather more information, our work goes further by utilizing an exploration-exploitation mechanism provided by Bayesian optimization to maximize human observations in the areas of interest. In contrast to the work of Molina et al. (2022), we demonstrate how robot exploration can be improved by correcting any systematic bias produced by robot's sensory system. Our main work contribution can complement any robot exploration technique by replacing 'Exploration Planner' module in Fig. 2 with any other exploration technique (e.g. entropy-based exploration).

3 Fully observable Poisson process

A *fully observable* Poisson process (FOPP) is a counting process $N(t)$ where a counter tells, without error, the number of events that occurred during a specified interval $[0, t)$. $N(t) = c_i$ states that in the i -th observation of interval $[0, t)$, there are c_i events. The number of events $N(t)$ in a finite interval of length t obeys the Poisson distribution

$$Poi(N(t) = c | \lambda) = \frac{e^{-\lambda} \lambda^c}{c!} \quad (1)$$

where λ represents the *arrival rate* in a fixed interval $[0, t)$. Bayesian estimation for fully observable Poisson processes is straightforward. Given a Gamma density

$$Gam(\lambda | \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda} \quad (2)$$

as a prior distribution over the parameter λ , where α, β are the shape and the rate parameters, the posterior over λ for a FOPP can be calculated via Bayesian inference with

$$\begin{aligned} P(\lambda | c_1, \dots, c_n) &\propto Poi(c_1, \dots, c_n | \lambda) Gam(\lambda | \alpha, \beta) \\ &= Gam\left(\lambda \mid \alpha + \sum_{i=1}^n c_i, \beta + n\right) \end{aligned} \quad (3)$$

This adds the sample counts $\sum_{i=1}^n c_i$ to the hyper-parameter α of the gamma prior, and adds the number of observations n to the hyper-parameter β of the gamma prior.

The FOPP model requires a single reliable sensor. With an unreliable sensor, FOPP inferences will be incorrect.

4 The partially observable Poisson process

The *partially observable* Poisson process (POPP) is a counting process $N(t)$ with arrival rate λ where the number of events appearing over the time interval $[0, t)$ is observed by one or more *unreliable* counters. The concept was firstly introduced in Jovan et al. (2018). The definition brings a distinction between the *true count* (or simply *count*), which refers to the number of events that actually occurred, and the *sensed count*, which refers to the count obtained by a counter (or sensor). Let c_i represent the true count over the interval $[0, t)$ during the i -th observation. With m counters unreliably observing c_i , we use $s_{j,i}$ to represent the sensed count given by sensor j in the i -th observation within the interval $[0, t)$ with $1 \leq j \leq m$. Let $\mathbf{s}_i = (s_{1,i}, \dots, s_{m,i})$ represent a vector of sensed counts from m sensors for the i -th observation of the process.

Figure 3 presents the graphical model derived from the definition of the POPP. This shows that the true count c_i has become a latent variable which can only be inferred from the sensed count. The posterior of λ is then inferred from the posterior of c_i after n observations, $i = 1 \dots n$.

The rate parameter λ of the POPP model can be inferred by marginalising over all possible true count values c_i and in the distribution of true counts given sensed counts $P(c_i | \mathbf{s}_i)$. Given n observations of the underlying process, let all




Fig. 3 Graphical representation of the partially observable Poisson process. Observable random variables are represented in shaded grey

observed true counts be represented by $\mathbf{c} = (c_1, \dots, c_n)$, and all sensed counts by $\mathbf{s} = (\mathbf{s}_1 \dots \mathbf{s}_n)$, for $1 \leq i \leq n$ (recalling each \mathbf{s}_n is produced by m sensors). The posterior of λ is then:

$$P(\lambda | \mathbf{s}) = \sum_{c_n=0}^{\infty} \dots \sum_{c_1=0}^{\infty} P(\lambda | \mathbf{c}) P(\mathbf{c} | \mathbf{s}) \quad (4)$$

where true count probabilities, $P(\lambda | \mathbf{c})$, can be drawn from the original FOPP definition:

$$P(\lambda | \mathbf{c}) = \text{Gam}\left(\lambda \mid \alpha + \sum_{i=1}^n c_i, n + \beta\right) \quad (5)$$

If we assume that the sensor counts for observation period i are conditionally independent (i.e. *uncorrelated*) given the true count c_i , then the probability of a collection of observations given the true count is defined as follows:

$$P(\mathbf{s}_i | c_i) = \prod_{j=1}^m P(s_{j,i} | c_i) \quad (6)$$

Using this, the probability of a particular sequence of n counts, given a sequence of n observations each from m sensors, $P(\mathbf{c} | \mathbf{s})$, can be defined as:

$$\begin{aligned} P(\mathbf{c} | \mathbf{s}) &\propto P(\mathbf{s}_1, \dots, \mathbf{s}_n | \mathbf{c}) P(\mathbf{c}) \\ &\propto \prod_{i=1}^n P(\mathbf{s}_i | c_i) P(c_i | c_{i-1}, \dots, c_1) \\ &\propto \prod_{i=1}^n \prod_{j=1}^m P(s_{j,i} | c_i) P(c_i | \mathbf{c}_{-1}) \end{aligned} \quad (7)$$

where $\mathbf{c}_{-1} = c_{i-1}, \dots, c_1$ ¹ and $P(c_i | \mathbf{c}_{-1})$ can be calculated as:

¹ \mathbf{c}_{-1} does not exist whenever $i = 1$, and $P(c_i) = \int_{\lambda=0}^{\infty} P(c_i | \lambda) \text{Gam}(\lambda | \alpha, \beta) d\lambda$

$$P(c_i | \mathbf{c}_{-1}) = \int_{\lambda=0}^{\infty} P(c_i | \lambda) P(\lambda | \mathbf{c}_{-1}) d\lambda \quad (8)$$

To complete Eq. 7 we must also define $P(s_{j,i} | c_i)$. The Poisson limit theorem states that the Poisson distribution may be used as an approximation to the binomial distribution (Papoulis & Pillai, 2002). Using this theorem as the foundation, an arbitrarily close approximation to the probability $P(s_{j,i} | c_i)$ is defined by assuming there exists a small enough finite subinterval of length δ for which the probability of more than one event occurring is less than some small value ϵ and that δ is small enough that ϵ is negligible. With this assumption, interval $[0, t)$ is split into l smaller subintervals I_1, \dots, I_l of equal size, with the condition that $l > \lambda$. Consequently, the whole interval $[0, t) = I_1, \dots, I_l$ becomes a series of Bernoulli trials, where the k^{th} trial corresponds to whether (1) an event e_k happens with probability λ/l and (2) a sensor j captures the event e_k as the detection d_k at the subinterval I_k .

Following this, $P(s_{j,i} | c_i)$ can be defined using of the count of true positives given c_i subintervals, and the false positives given the remaining $l - c_i$ subintervals. Let the probability of a *true positive detection* (TP) for sensor j in a single subinterval be $\tau_j = P_j(d | e=1)$, and the probability of a *false positive detection* (FP) be $\xi_j = P_j(d | e=0)$. Thus $P(s_{j,i} | c_i)$ is defined as a sum over all possible sensed counts of the product of two binomial distributions $B(r | n, \pi)$:

$$\begin{aligned} P(s_{j,i} | c_i) &= \sum_{r=0}^{c_i} B(r | c_i, \tau_j) \\ &\quad B((s_{j,i} - r) | (l - c_i), \xi_j) \end{aligned} \quad (9)$$

where the first binomial provides the probability of getting some proportion of the count from TP detections and the second binomial provides the probability of getting the remainder from FP detections.

Equation 4 shows the difficulty of estimation in the POPP model. Since no conjugate density provides an analytical solution for the posterior over λ , every sensed count \mathbf{s}_i must be retained to calculate the posterior of λ . That means elements representing each value of c_i on each observation grow infinitely. Even with an upper bound l on the maximum value of c_i , the number of elements to retain on each observation periods grows exponentially.

λ estimators

To address this difficulty, in (Jovan et al., 2018) we proposed three estimators, each of which offers an approximation to the true posterior $P(\lambda | \mathbf{s})$. The estimators are: (1) a *gamma filter*, which approximates Eq. 4 with a single gamma distribution minimising the KL-divergence $D_{KL}(P(\lambda | \mathbf{s}) || \text{Gam}(\lambda |$




Fig. 4 Average KL-divergence from the gamma and switching filters to $P(\lambda | s)$. The horizontal axis shows the true positive rate (top) and true negative rate (bottom) of one simulated sensor. The figure is taken from Jovan et al. (2018)

α, β) by gradient descent. The accuracy of this filter deteriorates as sensor reliability degrades. However, computation time is constant on each observation and Eq. 8 has a closed form, using the negative binomial distribution

$$\begin{aligned}
 P(c_i | \mathbf{c}_{-1}) &= \int_{\lambda=0}^{\infty} P(c_i | \lambda) P(\lambda | \mathbf{c}_{-1}) d\lambda \\
 &= \int_{\lambda=0}^{\infty} \text{Poi}(c_i | \lambda) \text{Gam}(\lambda | \alpha_{-1}, \beta_{-1}) d\lambda \\
 &= NB\left(c_i \mid \alpha + \sum_{k=1}^{i-1} c_k, \frac{\beta + (i-1)}{\beta + i}\right).
 \end{aligned} \tag{10}$$

with the hyperparameters $\alpha_{-1} = \alpha + \sum_{k=1}^{i-1} c_k$, $\beta_{-1} = \beta + (i-1)$ in $\text{Gam}(\lambda | \alpha_{-1}, \beta_{-1})$ being the updated prior distribution of λ up to \mathbf{c}_{-1} ;

(2) a *histogram filter*, which approximates Eq. 4 with a discrete distribution $Q(\lambda | s)$ by quantising λ . The advantage of this filter over the gamma filter is that it can track the posterior to an arbitrary fidelity via a finer quantisation with the cost of computation time. Its disadvantage is an increase in computation time compared to the gamma filter;

(3) a *switching filter*, which approximates Eq. 4 either by a gamma filter or by a histogram filter depending on whether $P(\lambda | s)$ resembles a gamma distribution and can be approximated by the gamma filter via KL-divergence $D_{KL}(P(\lambda | s) || \text{Gam}(\lambda | \alpha, \beta))$.

In general (and in our experimental work from Sect. 6 onwards) we use the switching filter as the estimator to the true posterior $P(\lambda | s)$ because it combines the best of both the gamma filter (fast calculation) and the histogram filter (accurate approximation) with minimum loss in similarity to the true posterior $P(\lambda | s)$. Figure 4 shows KL-divergence between the gamma and switching filters to the true posterior over different sensor reliabilities using simulated data. Note

that the histogram filter was not included because it perfectly tracked $P(\lambda | s)$, i.e., $D_{KL}(P(\lambda | s) || Q(\lambda | s)) \approx 0$. A more detailed presentation of these estimators is given in (Jovan et al., 2018).

5 The POPP extensions

In (Jovan et al., 2018) we demonstrated that the POPP model is able to efficiently correct miscounts made by multiple unreliable counting devices observing a single Poisson process. However, the POPP model is limited by two assumptions:

- (1) the sensors are conditionally independent given the true count, and
- (2) the degree of the unreliability of a sensor (i.e. τ and ξ) is precisely known.

In this paper, we propose three extensions to the POPP model to tackle these assumptions. The first extension (POPP-Beta) extends the POPP model with an observation model which captures uncertainty about the role of the sensor reliability. The second extension (C-POPP) modifies the POPP model to accommodate correlations between sensors. The third extension (POPP-Dirichlet) combines these ideas to jointly address both assumptions.

5.1 POPP-Beta

The POPP model requires the true positive and false positive rates to be specified for sensor j , i.e. $\tau_j = P_j(d | e=1)$ and $\xi_j = P_j(d | e=0)$. The POPP model requires these rates to be accurate in order to generate correct posteriors over λ . To accurately determine the rates in practice, one needs to have a large data set of both sensed counts and the ground truth. Given the ground truth is typically manually created, this places a large burden on experts who need to label the data.

Here, we extend the original POPP model to take into account uncertainty in the true and false positive rates due to limited training data. To model this uncertainty we use Bayesian estimation to determine the true positive rate (τ) and false positive rate (ξ). We use Beta distributions as priors for τ and ξ because the Beta distribution act as a conjugate to the binomial distribution, providing a family of prior probability distributions for the parameter of a binomial distribution. The Beta-binomial conjugacy leads to an analytically tractable compound distribution called the Beta-binomial distribution $BB(d | c, \zeta, \eta)$, where the p parameter in the binomial distribution $B(d | c, p)$ is drawn from a Beta distribution $Be(p | \zeta, \eta)$.

Our sensor rates, τ_j and ξ_j , are now estimated from two Beta distributions: $Be(\tau | \zeta_\tau, \eta_\tau)$ and $Be(\xi | \zeta_\xi, \eta_\xi)$. ζ_τ and




Fig. 5 Graphical representation of the POPP-Beta. Instead of having fixed estimated points for the sensor rates τ and ξ like in the POPP model, they are represented by Beta distributions in the POPP-Beta

ζ_ξ are the number of true positive and false positive detections in the ground truth data respectively. η_τ and η_ξ are the number of true negative and false negative detections in the ground truth data respectively. Given these parameters, we form the POPP-Beta model from POPP by replacing Eq. 9 with:

$$P(s_{j,i} | c_i) = \sum_{r=0}^{c_i} BB(r | c_i, \zeta_\tau, \eta_\tau) BB(\delta_s r | \delta_c r, \zeta_\xi, \eta_\xi) \quad (11)$$

with $\delta_s r = (s_{j,i} - r)$, and $\delta_c r = (l - c_i)$.

With a sensor model which follows beta densities and is fully integrated, as a distribution, in the sensed count likelihood $P(s_{j,i} | c_i)$ as shown in Equation 11, we obtain a graphical model with the structure shown in Fig. 5. One should note that the difference between the POPP and POPP-Beta model, lies only in the change from Eq. 9 to 11. However, given little training data for the observation model, the POPP-Beta model is expected to be more conservative in estimating the posterior $P(\lambda | \mathbf{s})$ over λ than the POPP model.

5.2 Correlated POPP

Recall that Eq. 6 is defined under the assumption that each sensor count is conditionally independent from all the others given the true count. This assumption ignores the correlations between sensors. To introduce correlations between sensors we must alter Eq. 6 and Eq. 9 from the POPP model.

Recall that the probability of a particular sensed count given the true count $P(\mathbf{s}_i | c_i)$ was defined from the Poisson limit theorem as a sequence of Bernoulli trials over l subintervals. With correlated sensors, the observation of an event e_k in the k^{th} trial no longer follows the Bernoulli distribution. Instead it follows the categorical distribution, where the

k^{th} trial corresponds to whether a particular combination of binary detections $d_{1,k}, \dots, d_{m,k}$ happens in subinterval I_k . Therefore, we move our notation from using $s_{j,i}$ representing sensed counts for particular sensor j independently at time interval i to a matrix representing m sensor detections together at time interval i . Formally, we replace Eq. 6 and Eq. 9 with a probability of a series of detection outcomes given the true count c_i at interval i as the following.

We first define for some interval i , l subintervals, and m sensors, there is a binary matrix of detections \mathbf{D}^i ². $\mathbf{D} \in \mathcal{D}^{m,l}$ the set of binary matrices of dimension $m \times l$. Each column k of \mathbf{D} , we denote $\mathbf{D}_{:,k} = \mathbf{d} = \{0, 1\}^m$ with $k = 1, \dots, l$. $\mathbf{D}_{:,k}$ is a vector of detections from m different sensors at particular subinterval k .

We further define $e_k \in \{0, 1\}$ as the variable indicating whether or not an event is hypothesized to have occurred in sub-interval k . $e_k = 1$ means that an event occurred. We define P^+ as the categorical distribution of \mathbf{d} , conditioned on $e = 1$, i.e.

$$P^+(\mathbf{d}) = P(\mathbf{d} | e = 1) \quad \forall \mathbf{d} \in \{0, 1\}^m \quad (12)$$

and, by analogy,

$$P^-(\mathbf{d}) = P(\mathbf{d} | e = 0) \quad \forall \mathbf{d} \in \{0, 1\}^m \quad (13)$$

Both P^+ and P^- have 2^m elements³. These two probabilities represent *true positive rates* and *true negative rates* as τ and ξ for the POPP model. Similar to τ and ξ , P^+ and P^- are estimated from both detections of each sensor and the corresponding actual (non-)event as ground truth. However, unlike τ and ξ which are sensor specific, P^+ and P^- consider all combinations of binary detections from sensors given the true event. This means the number of elements in P^+ and P^- grows by a factor of two for each sensor added. Due to the size of P^+ and P^- , they may need more than a few hundred of detections together with their corresponding events to be estimated.


We can partition the subintervals $1, \dots, l$ into two sets. \mathbf{e}^+ is the set of subintervals k where $e_k = 1$, and \mathbf{e}^- is the set of subintervals k where $e_k = 0$. We can define a partition of the subintervals by a pair $(\mathbf{e}^+, \mathbf{e}^-)$. The set of possible partitions such that \mathbf{e}^+ has a fixed size c , i.e. $|\mathbf{e}^+| = c$, is denoted Σ_c , so that $(\mathbf{e}^+, \mathbf{e}^-) \in \Sigma_c$.

We further define $\mathbf{D}_{\mathbf{e}^+}$ as an $m \times c$ detection matrix formed from all the columns $\mathbf{D}_{:,k}$ where $k \in \mathbf{e}^+$, and $\mathbf{D}_{\mathbf{e}^-}$ as the corresponding $m \times (l - c)$ detection matrix formed from all the columns $\mathbf{D}_{:,k}$ where $k \in \mathbf{e}^-$.

² We drop the (i) for all notations in this subsection as we will consider a single interval

³ \mathbf{d} is dropped in the representation unless the context is not clear.

Fig. 6 Graphical representation of C-POPP. Unlike the POPP model, the matrix detection \mathbf{D} represents a joint detection at particular time interval and is affected by the value of the true count c , and the sensor rates (joint true positive rate P^+ and joint true negative rate P^-)



As there may be duplicate columns in either or both \mathbf{D}_{e^+} and \mathbf{D}_{e^-} , we define a count vector for each.

$$\mathbf{g}^+ = \text{count}(\mathbf{D}_{e^+})$$

and

$$\mathbf{g}^- = \text{count}(\mathbf{D}_{e^-})$$

such that $\sum_{q=1}^{2^m} g_q^+ + \sum_{r=1}^{2^m} g_r^- = l$ where each of $\mathbf{g}^+, \mathbf{g}^-$ are of length 2^m , having one element for every possible detection vector $\mathbf{d} \in \{0, 1\}^m$.

In order to define the joint probability of a particular count being yielded by a particular sequence of detection outcomes, we must consider all possible combinations of true positives and false positives that could be generated by that sequence by exploring all elements of Σ_c . We do this in the following definition of $P(\mathbf{D} | c)$, and define the probability of a given sequence of detection groups yielding count c using the multinomial distribution.

$$P(\mathbf{D} | c) = \sum_{(\mathbf{e}^+, \mathbf{e}^-) \in \Sigma_c} M(\mathbf{g}^+ | c, P^+) M(\mathbf{g}^- | \delta_l c, P^-) \quad (14)$$

with $\delta_l c = (l - c)$.

Equation 14 can be understood by analogy to Eq. 9. In both equations all possible ways pairs of true and false positives counts which sum to c are considered. In the conditionally independent case the binomial distribution is used to determine the probability of each count from the available trials given the true and false positive rates. However, in the conditionally independent case, Eq. 9 is calculated independently for each sensor, and the joint probability of those sensors results in Eq. 6. In the correlated case the multinomial distribution is used to determine the probability of each count from a possible sequence of joint observations and their probability of yielding a count. With that, Eq. 14 removes the need of Eq. 6 in C-POPP model. A graphical representation for C-POPP can be seen in Fig. 6.

One should note that the benefit of C-POPP is that it exploits correlations among multiple sensors contributing to detection counts. If there is only one sensor counting events, then C-POPP collapses to POPP.

5.3 The POPP-Dirichlet

The C-POPP model requires the true positive rate P^+ and true negative rate P^- to be specified in advanced in estimating the parameter λ of a Poisson process. These are an extension of τ and ξ where the rates provide a probability for a particular combination of binary detections coming from each sensor given the true event as shown in Eq. 12 and Eq. 13.

To construct an observation model of P^+ and P^- , one needs to have both detections and the corresponding actual (non-)events as ground truth. Pre-processing involving expert interventions is typically required before the detections and their corresponding ground truth can be further used. Similarly to the POPP model, the C-POPP model requires the observation model to be accurate to avoid the posterior over λ drifting away from the true posterior. If attaining an accurate observation model for the POPP model is a problem, then this becomes more challenging in the case of C-POPP model. This is because the training data needed to construct an observation model grows by a factor of two for each sensor involved.

Analogously to the extension from the POPP model to the POPP-Beta model, we can expand the C-POPP observation model. In this case the observation models (P^+ and P^-) will follow Dirichlet distributions. The Dirichlet distribution is an appropriate distribution since P^+ and P^- are the probabilities of categorical distributions which set the probabilities of multinomial distributions in Eq. 14 and Dirichlet distributions provide a family of conjugate prior probability distributions for the multinomial distribution. The Dirichlet-multinomial conjugacy leads to an analytically tractable compound distribution which is called the Dirichlet-multinomial distribution, where the $\mathbf{p} = (p_1, \dots, p_r)$ parameter in the multinomial distribution $Mult(\mathbf{d} | c, \mathbf{p})$ is randomly drawn from a Dirichlet distribution $Dir(\mathbf{p} | \boldsymbol{\zeta})$.

$$\begin{aligned} P(\mathbf{d} | c, \boldsymbol{\zeta}) &= \int P(\mathbf{d} | c, \mathbf{p}) P(\mathbf{p} | \boldsymbol{\zeta}) d\mathbb{S}_r \\ &= \int Mult(\mathbf{d} | c, \mathbf{p}) Dir(\mathbf{p} | \boldsymbol{\zeta}) d\mathbb{S}_r \\ &= DM((d_1, \dots, d_r) | c, (\zeta_1, \dots, \zeta_r)) \end{aligned} \quad (15)$$

with $\mathbf{d} = (d_1, \dots, d_r)$, $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_r)$, and $d\mathbb{S}_r$ denotes integrating \mathbf{p} with respect to the $(r - 1)$ simplex⁴.

⁴ The support of the Dirichlet distribution is the $(r - 1)$ -dimensional simplex \mathbb{S}_r ; that is, all r dimensional vectors which form a valid probability distribution




Fig. 7 Graphical representation of POPP-Dirichlet. Instead of having fixed estimated points for the (joint) sensor rates P^+ and P^- , they are represented by Dirichlet distributions in the POPP-Dirichlet

Given m sensors, an observation model is now represented as two Dirichlet distributions: $Dir(P^+ | \zeta^+)$, and $Dir(P^- | \zeta^-)$ with $\zeta^+ = (\zeta_0^+, \dots, \zeta_{(m^2)-1}^+)$ and $\zeta^- = (\zeta_0^-, \dots, \zeta_{(m^2)-1}^-)$. ζ^+ and ζ^- set the overall shape of the Dirichlet priors, with each ζ_q term counting the number of times that particular combination of sensor detections were produced given a positive ($\zeta^+, e = 1$) or negative ($\zeta^-, e = 0$) detection.

Given a joint sensor model where its elements follow a Dirichlet density and several Dirichlet-multinomial distributions, which provide an unconditional distribution of \mathbf{d} , we replace Eq. 14 with:

$$P(\mathbf{D} | c) = \sum_{(e^+, e^-) \in \Sigma_c} DM(\mathbf{g}^+ | c, \zeta^+) DM(\mathbf{g}^- | \Delta, \zeta^-) \quad (16)$$

with $\Delta = (l - c)$, Σ_c and \mathbf{D} as defined in Sect. 5.2.

The difference between the C-POPP model and the POPP-Dirichlet lies only in Eq. 14 being replaced by 16 which is depicted by Fig. 7. The difference makes the POPP-Dirichlet to be more conservative in estimating the posterior $P(\lambda | \mathbf{s})$ over λ than the C-POPP model given a certain Dirichlet prior, and limited training data for the sensor model.

6 Evaluation on synthetic data

In this section we evaluate POPP and its extensions on synthetic data to demonstrate the properties of these models when estimating the arrival rate λ of a Poisson process. With synthetic data, sensor reliability can be controlled, and the true λ and the true counts c_i can be known for each sample.

In our experiments we initially generate a training set of $n = 12$ (true) counts from a Poisson process $P(c | \lambda' = 3)$ with a time interval $t = 10$ time unit. Along with the training

set count c_1, \dots, c_{12} , for each count c_i , we also generate the corresponding event occurrence $e_k \in \{0, 1\}$ on each subinterval $k \in \{1, \dots, t\}$, a sensed count \mathbf{s}_i , and \mathbf{D}_i from two unreliable sensors with 10 subintervals for each sensed count \mathbf{D}_i (i.e. $m = 2, l = 10, \mathbf{D} \in \mathcal{D}^{2,10}$ in our evaluation). To capture a range of possible sensor correlations and performance characteristics, the sensed counts for the training set are produced from 12 different sensor configurations (see Table 1). The true and sensed counts are then used to build (joint where appropriate) sensor models for the POPP extensions described above. For the POPP-Beta and the POPP-Dirichlet models, we set the hyperparameters of the Dirichlet prior and Beta prior to follow uniform distribution, i.e., $\zeta_\tau = \eta_\tau = \zeta_\xi = \eta_\xi = 1$ for the POPP-Beta, and $\zeta^+ = \zeta^- = (1, 1, 1, 1)$ for the POPP-Dirichlet. Most of these (hyper) parameters ($l, t, \zeta, \eta, \zeta^+, \zeta^-$), except the number of sensors m , are reused in our real-world experiment in the next chapter.

We then generate a new set of $n = 144$ true counts and the corresponding sensed counts for each of the 12 sensor configurations. These sensing are used as input in a filtering process to estimate the posterior of λ according to each of the four models defined above (POPP, POPP-Beta, C-POPP, and POPP-Dirichlet), plus FOPP. We chose the training set size $n = 12$ such that there is insufficient data to build an accurate sensor model. This allows the POPP-Dirichlet and the POPP-Beta models to compensate with loose Dirichlet and beta densities.

The 12 sensor configurations mentioned previously represent 12 different experimental conditions under which we can test our proposed models. In six of the configurations we vary the true joint positive rates (true P^+) of the two sensors whilst fixing their true joint negative rates (true P^-). In the other six we fix the true joint positive rates (TJPRs) whilst varying the true joint negative rates (TJNRs). Both cases cover variations where the sensors are uncorrelated, positively correlated and negatively correlated, and in each case where the overall true (positive or negative) rates are either high (0.9) or low (0.1). The detailed configurations are presented in Table 1.

The performance of all POPP models was assessed by measuring how accurate each model is in estimating the true λ' . The true λ' is estimated by applying the FOPP model on the true counts⁵. Two options were used to measure the accuracy: (1) the RMSE of the expectation (mean) and the MAP hypothesis (mode) of each model posterior distribution over λ to the true λ' ; and (2) the Jensen-Shannon distance between the posterior distribution $P(\lambda | \mathbf{s}_i)$ and the distribution of the true λ' .

⁵ This FOPP model is not the one being tested with the POPP model. This FOPP model is built using the true count data, while the FOPP model being tested is built using the noisy sensor data.

Table 1 The sensor configurations for the evaluation on synthetic data. “+ corr” and “- corr” mean a positive correlation and a negative correlation between two sensors respectively

	e_k $d_{1,k}, d_{2,k}$	1				0			
		0, 0	0, 1	1, 0	1, 1	0, 0	0, 1	1, 0	1, 1
TJPR	TJNR								
Low + corr	Fixed	0.1	0.0	0.0	0.9	1.0	0.0	0.0	0.0
High + corr	Fixed	0.9	0.0	0.0	0.1	1.0	0.0	0.0	0.0
Low - corr	Fixed	0.0	0.05	0.05	0.9	1.0	0.0	0.0	0.0
High - corr	Fixed	0.0	0.45	0.45	0.1	1.0	0.0	0.0	0.0
No correlation	Fixed	0.033	0.033	0.033	0.901	1.0	0.0	0.0	0.0
No correlation	Fixed	0.3	0.3	0.3	0.1	1.0	0.0	0.0	0.0
Fixed	Low + corr	0.0	0.0	0.0	1.0	0.9	0.0	0.0	0.1
Fixed	High + corr	0.0	0.0	0.0	1.0	0.1	0.0	0.0	0.9
Fixed	Low - corr	0.0	0.0	0.0	1.0	0.9	0.05	0.05	0.0
Fixed	High - corr	0.0	0.0	0.0	1.0	0.1	0.45	0.45	0.0
Fixed	No correlation	0.0	0.0	0.0	1.0	0.9	0.033	0.033	0.033
Fixed	No correlation	0.0	0.0	0.0	1.0	0.1	0.3	0.3	0.3




Fig. 8 The RMSE of posterior estimates of λ for the POPP and its variation models with 12 sample data used to build the (joint) sensor model with variation in \mathcal{P}^+ . All models are compared to the FOPP model. Each trial consisted of a stream of $\mathbf{s}_1 \dots \mathbf{s}_{144}$ samples to update $P(\lambda | \mathbf{s}_i)$. Accuracies of MAP estimates are shown in the top panel, accuracies of expectation of the posterior in the bottom panel. Each data point is an average of 30 trials. Standard errors are shown




Fig. 9 The Jensen-Shannon distance of posterior estimates of λ for the POPP and its variation models with 12 sample data used to build the (joint) sensor model with variation on \mathcal{P}^+ . All models are compared to the FOPP model. Each trial consisted of a stream of $\mathbf{s}_1 \dots \mathbf{s}_{144}$ samples to update $P_G(\lambda | \mathbf{s}_i)$. Each data point is an average of 30 trials. Standard errors are shown




Fig. 10 The RMSE of posterior estimates of λ for the POPP and its variation models with 12 sample data used to build the (joint) sensor model with variation in \mathcal{P}^- . All models are compared to the FOPP model. Each trial consisted of a stream of $\mathbf{s}_1 \dots \mathbf{s}_{144}$ samples to update $P(\lambda | \mathbf{s}_i)$. Accuracies of MAP estimates are in the top panel, accuracies of the expectation of the posterior in the bottom panel. Each data point is an average of 30 trials. Standard errors are shown




Fig. 11 The Jensen-Shannon distance of posterior estimates of λ for the POPP and its variation models with 12 sample data used to build the (joint) sensor model with variation on \mathcal{P}^- . All models are compared to the FOPP model. Each trial consisted of a stream of $\mathbf{s}_1 \dots \mathbf{s}_{144}$ samples to update $P_G(\lambda | \mathbf{s}_i)$. Each data point is an average of 30 trials. Standard errors are shown

Figures 8 and 9 show the accuracy of all POPP models over the variations of TJPR (P^+), whereas Figs. 10 and 11 show the accuracy across the variation in TJNR (P^-). From these figures, POPP-Beta and POPP-Dirichlet show a better accuracy than POPP and C-POPP. C-POPP and POPP-Dirichlet, which utilize correlations among sensors to estimate the arrival rate λ' , tend to be more accurate than the standard POPP and POPP-Beta. In general, POPP-Dirichlet tends to be more accurate than any other POPP model thanks to its ability to model correlation among sensors *and* how confident it is in its sensor model. One should note that if the number of training samples for the (joint) sensor model is high, then the POPP-Dirichlet and the C-POPP should have similar posterior distributions. This is because the POPP-Dirichlet will have tight densities over the sensor models, and these should be comparable to the point estimates used in the C-POPP sensor models.

In this paper, we remove computation time per sample analysis between POPP and its extensions because the computation relies heavily on the filters chosen. The time to calculate the distribution of sensed count given the actual count between POPP, POPP-Beta, C-POPP and the POPP-Dirichlet on each sample can be considered constant and, therefore, is negligible to the total computation time. Our prior work provided a detailed comparison in computational efficiency between different filters (Jovan et al., 2018).

7 Evaluation on aggregate human occupancy behaviour dataset

We now investigate the performance of the POPP model and its extensions on a real world dataset⁶. The dataset was gathered from an office building in which a mobile robot (Hawes et al., 2016) counted the number of people in different regions whilst patrolling (see Fig. 12 for the map of the building). The dataset contains time series counts from three different automated person detectors (Dondrup et al., 2015). These use laser, depth camera and RGB information. We refer to these detectors respectively as the leg detector (LD), upper body detector (UBD), and change (or scenery) detector (CD). Each of these detectors acts as one sensor. Each returns a sensed count of the number of people it detected in each 10 minute interval during the day. To unify different frequency of detections of each sensor, we used the lowest frequency detection from the change detector and limited to maximum one detection per minute. These detectors are unreliable, as can be seen from Fig. 13, which shows examples of correct and incorrect detections.

⁶ The dataset can be downloaded from https://github.com/ferdianjovan/spectral_popp.



Fig. 12 The office building in which the robot gathered data. Areas are bounded by imaginary lines. The figure is taken from Jovan et al. (2018)




Fig. 13 Correct and incorrect detections, and non-detections, from different regions in the environment for each sensor. Top row: change detector. Middle row: upper body detector. Bottom row: leg detector. The figure is taken from Jovan et al. (2018)

By comparing the ground truth with the detections made by sensors, we compute a sensor model for each region. An average of the sensor models across all regions can be seen in Table 2. Although the robot operated for 24 hours day, the sensor models were built using only the data collected from 10am to 8pm, since there were few detections outside these times. From a 69 day trial of the mobile robot, we obtained 48 days of usable observations. We specified a time interval for each Poisson distribution of 10 minutes, and recorded both the true counts and the detections made by each sensor in each interval. We assumed the underlying process in each region to be a periodic Poisson process in which there is a one-day periodicity, i.e. $\lambda(t) = \lambda(t + \Delta)$ with $\Delta = 24 \times 60$ (minutes). This means that the expected number of people each day at a particular time is expected to be the same across the 48 days of observations. We estimated the true parameter $\lambda'(t)$ of the Poisson distribution at t by running a FOPP model on the true counts within each interval. We use this estimate of $\lambda'(t)$ from the true counts as the target which the POPP models must estimate from the sensed counts.

The different POPP approaches rely on sensor models that must be calculated from a confusion matrix relating true counts to the sensed counts from the different sensors. To

Table 2 Averaged sensor models across all areas trained from 48 days of data

Sensor	True positive	True negative
Leg detector	0.387	0.951
Upper body detector	0.356	0.882
Change detector	0.731	0.900




separate the training and testing data we performed four fold cross-validation with data splits being on whole days, i.e., we used 12 days of data as a training set for a sensor model and then used the remaining 36 days of data as a test set on which to test the inferences made by each model from the sensor counts.

For the 36 days of test data, the different models each made predictions of the $\lambda(t)$ parameter of the Poisson. Given this, we recorded (1) the RMSE between the MAP hypothesis of each model posterior distribution over $\lambda(t)$ and the true $\lambda'(t)$ and (2) the Jensen-Shannon distance between the posterior distribution $P(\lambda(t) | s_i)$ and the distribution of the true $\lambda'(t)$. Using these metrics, we compared the performance of all POPP models (estimated using the switching filter described in Sect. 4) to the Bayes' filter arising from the FOPP model. The FOPP model is a single sensor model and was estimated from the change detector counts since this was the most reliable detector among the three available (as shown in Table 2).

Figures 14 and 15 show the accuracy comparison between all POPP models and the standard FOPP model over time. It can be seen that all models become more accurate as the days pass. All POPP models show more accuracy over the standard FOPP model. The $\lambda(t)$ estimate produced by the POPP-Dirichlet model is more accurate than the ones produced by the standard POPP model and the POPP-Beta model. However, the POPP-Dirichlet estimate is not always more accurate than the one produced by the C-POPP model.

As the POPP-Dirichlet model is more conservative in estimating the parameter $\lambda(t)$ than the C-POPP model, the estimate moves more slowly towards the true $\lambda'(t)$. This is seen in Fig. 15. By the third day, the POPP-Dirichlet model outperformed the POPP, POPP-Beta, and C-POPP models in terms of accuracy. However, the accuracy gap between the C-POPP model and the POPP-Dirichlet model becomes smaller over time. By the 36th day the C-POPP model outperforms the POPP-Dirichlet by a small margin. It should be noted that Figs. 14 and 15 are averaged RMSE and the Jensen-Shannon distance from 10 different regions over time. The more regions with high volume of data available, the more accurate the joint sensor model, especially for C-POPP, will be and, in turn, the more accurate the C-POPP filter becomes in estimating the parameter $\lambda(t)$.

Figures 16 and 17 show the RMSE and Jensen-Shannon comparison between all POPP models and the FOPP across

**Fig. 14** The RMSE evolution of periodic Poisson processes with POPP, POPP-Beta, C-POPP, POPP-Dirichlet and FOPP filters from day 3 to day 36, averaged across all regions. Standard error is shown**Fig. 15** The Jensen-Shannon distance evolution of the FOPP, the POPP, the POPP-Beta, the C-POPP, and the POPP-Dirichlet filters in periodic Poisson processes from day 3 to day 36 in a 3-day interval, averaged across all regions. Standard error is shown**Fig. 16** The RMSE of the FOPP, POPP, POPP-Beta, C-POPP, and POPP-Dirichlet filters across regions. The RMSE(s) are taken at the 36th day. Standard error is shown

different regions by the end of the 36th day. It can be seen that the POPP-Dirichlet and the C-POPP once again outperformed the other models. Some regions such as 1, 2, and 3 have much more data than other regions. Since this provides more data to create the sensor models than other regions, the point-estimate joint sensor model for the C-POPP filter can be more accurately estimated for these regions. Unlike C-POPP filter, the POPP-Dirichlet estimates the joint sensor model as a distribution. This drives the POPP-Dirichlet




Fig. 17 The Jensen-Shannon of the FOPP, POPP, POPP-Beta, C-POPP, and POPP-Dirichlet filters across regions. The Jensen-Shannon value(s) are taken at the 36th day. Standard error is shown

slower and more conservative in estimating the parameter $\lambda(t)$ than the C-POPP model. Together with the choice of Dirichlet prior that follows uniform distribution, the POPP-Dirichlet requires more data to accurately estimate its joint sensor model.

The POPP-Dirichlet has an advantage on regions with low volume of data such as region 4, 5, 6 and 7. As some of these data were used to construct the joint sensor model for both C-POPP and the POPP-Dirichlet, a small amount of data creates an inaccurate point-estimate joint sensor model, which is used by the C-POPP filter. This problem is handled appropriately on the POPP-Dirichlet with its distribution joint sensor model with the help of Dirichlet prior as explained in Sect. 5.3.

One interesting finding here is that there is small to no difference in performance between the POPP and the POPP-Beta filters on region 4, 5, 6, and 7. One would have thought that the performance of these two filters should follow the C-POPP and the POPP-Dirichlet filters. We argue that the volume of data used to create the sensor models for both POPP and the POPP-Beta were enough for an accurate estimate of point-estimate sensor model (POPP) and distribution sensor model (POPP-Beta). However, due to high correlations among sensors which were not captured by both the POPP and the POPP-Beta sensor models, the accuracy in estimating the parameter $\lambda(t)$ is worse than the C-POPP and the POPP-Dirichlet. It is also worse for the POPP-Beta filter since the POPP-Beta is more conservative in estimating the parameter $\lambda(t)$ than the POPP model. For example, region 4 contains high tables and tall chairs where the leg detector tended to falsely detected them as a person. Unless an upper body detector detects a person, the leg detector detection may be ignored. On the other hand, region 7 is a hallway with a water dispenser around the corner. This water dispenser is often falsely detected as a person by the upper body detector and the leg detector detections helps in reducing this mistake.

8 Exploring for human activities

So far, the paper has focused on Bayesian methods for inferring a belief state about the spatio-temporal patterns of human occupancy from unreliable sensors. Given such a belief state a robot may plan how to actively explore to acquire new information so as to complete a task (Hanheide et al., 2017; Sridharan et al., 2019). Here, the robot uses predicted counts from the belief state to *explore* so as to detect human activities with increasing efficiency.

Specifically, the robot's choice is whether to explore new region-time combinations or to exploit region-time combinations that are known to yield a high number of activities. This an instance of an *exploration-exploitation* problem. Exploration-exploitation problems arise whenever an agent lacks an adequate model of the process it must control. At each moment, the agent chooses either to explore so as to improve the model or to exploit the existing model so as to maximise immediate performance.

While exploration-exploitation problems in reinforcement learning, are typically intractable, there are well known, fast to compute, approximations (Wyatt, 1998; Alba & Dorrnsoro, 2005; Audibert et al., 2009). One such approach is to use the upper bound of a probability distribution over the quantity being maximised. This causes the decision-making agent to exploit high-scoring, certain estimates, and explore highly uncertain estimates. In our robot exploration, for example, when the robot visits a place, it can be because the place either actually has high number of people (*exploitation*) or potentially has high number of people (*exploration*). In our case we use an upper bound on the arrival rate (λ) of a Poisson process (λ_{UB}) to choose the region for the robot to visit next. The upper bound of the probability interval of the arrival rate of a Poisson process is calculated as follows:

$$\lambda_{UB}(t_i, t_j) = \int_{t_i}^{t_j} CDF^{-1}(\% = 0.95 \mid \alpha_t, \beta_t) dt \quad (17)$$

with $\lambda_{UB}(t_i, t_j)$ as the upper bound of λ within time t_i and t_j , $i, j \in \{1, \dots, \Delta\}$, and CDF^{-1} as the inverse of the cumulative density function of a Gamma distribution. Given the upper bounds $\lambda_{UB}^r(t_i, t_j)$ for each region r from the set of all regions R , the region to be visited between time t_i and t_j is chosen by:

$$\arg \max_{r \in R} \lambda_{UB}^r(t_i, t_j) \quad (18)$$

Figure 18 depicts a comparison between the MAP hypothesis estimate and the upper bound estimate of a Poisson process.

To tie the estimate of a particular Poisson process over a time interval to data collected previously, as in Sect. 7 we assume that human presence in each region follows a *periodic* Poisson process with daily periodicity. This allows us




Fig. 18 A spectral Poisson process of region 9 (see Fig. 12) represented by its MAP hypothesis (blue line) and its upper bound of the probability interval (red line) (Color figure online)

l-AAM Jovan et al. (2016)

Input: x_1, \dots, x_n : input signal,
total: maximum total frequency

Output: S : a collection of (s, p, f)

Procedure:

1. Init. $k \leftarrow 0$
// Get frequency 0 with Fourier Transform
2. $[s, p, f] \leftarrow DFT(x_1, \dots, x_n)[0]$
3. $S[k] \leftarrow [s, p, f]$
4. Repeat until $k > \text{total}$
 - $k \leftarrow k + 1$
// Get the highest amplitude's frequency
 - $[s, p, f] \leftarrow \arg \max_s DFT(x_1, \dots, x_n)$
// Update S with frequency f
 - if $f \in S, [s', p', f'] \leftarrow S[k', f' = f]$
 $s \leftarrow s + s'; p \leftarrow p + p'$
 - $S[k] \leftarrow [s, p, f]$
// Create a cosine signal from f
 - $x'_1, \dots, x'_n \leftarrow s * \cos(2\pi * f + p)$
// Subtract current x_1, \dots, x_n with the cosine
 - $x_1, \dots, x_n \leftarrow x_1, \dots, x_n - x'_1, \dots, x'_n$

to regularise, and fill missing data, across the point estimates of upper bounds using methods based on the Fourier transform. This exploits assumptions and algorithms introduced in our prior work. In particular, the series of upper bounds $\lambda_{UB}(t_i, t_j)$ are encoded and extracted via spectral analysis with the *l*-AAM technique described in (Jovan et al., 2016). The plot in Fig. 18 shows how a spectral Poisson process look like, i.e., the effects of the spectral processing on a periodic Poisson process. Algorithm 2 depicts the process of computing the upper bound of a Poisson process and applying spectral analysis to it. We use this approach with upper bounds produced by our previously presented estimators: FOPP, POPP, and POPP-Beta. C-POPP and POPP-Dirichlet estimators are excluded in our experiments due to a need to limit experimental time to 45 days to keep building use conditions that were broadly the same.⁷

⁷ The experiments were conducted during a single academic semester within a university building.

Algorithm 2 Upper Bound

Input: $(\alpha_1, \beta_1), \dots, (\alpha_n, \beta_n)$: Poisson process

Output: $\lambda_1^{ub}, \dots, \lambda_n^{ub}$: upper bound

Procedure:

1. Init. $k \leftarrow 1, m \leftarrow \eta$
2. Repeat until $k > n$
 - $k \leftarrow k + 1$
// Get the upper bound
 - $\lambda_k \leftarrow CDF(0.95, \alpha_k, \beta_k)$
- // Transform $\lambda_1, \dots, \lambda_n$ to with *l*-AAM
3. $S \leftarrow \mathbf{l-AAM}(\lambda_1, \dots, \lambda_n, m)$
5. Init. $k \leftarrow 0, \lambda_1^{ub}, \dots, \lambda_n^{ub} \leftarrow (0, \dots, 0)$
4. Repeat until $k > m$
 - // Create a cosine signal from $S[k]$
 - $[s, p, f] \leftarrow S[k]$
 - $x_1, \dots, x_n \leftarrow s * \cos(2\pi * f + p)$
// Add current $\lambda_1^{ub}, \dots, \lambda_n^{ub}$ with the cosine
 - $\lambda_1^{ub}, \dots, \lambda_n^{ub} \leftarrow \lambda_1^{ub}, \dots, \lambda_n^{ub} + x_1, \dots, x_n$

Exploration evaluation

The dataset used in the previous section was collected by a mobile robot over 69 days of a real world trial. This robot was controlled by the exploration models described above. Due to hardware failures, sensor malfunctions and other external issues, only 48 days from the dataset were usable.

Three different exploration models were applied separately during three phases of the 69 days of the trial. All of these models used Eq. 18 to create their exploration policies. For the first 27 day phase of the trial, the robot followed an exploration policy based on the FOPP model. This resulted in 18 days of data. From day 28 to day 47, the robot followed an exploration policy according to the POPP model. This resulted in 15 days of data. Finally, from day 48 onwards, the robot followed an exploration policy according to the POPP-Beta model. This also resulted in 15 days of data. Such that all three models can be compared equally, in the following we also constrain the data available for for the FOPP model to the first 15 of its 18 days. We can compare the different exploration policies on the observations the robot made during the phase each policy was active. Due to the absence of information regarding occupancy in the places that the robot did not visit, only a comparison of the positive observations can be made.

Figure 19 shows the percentage of visits to each region which yielded a non-zero true count. As can be seen, the exploration policy produced by POPP-Beta has the highest proportion of such visits in many of the regions, followed by the exploration policy according to the POPP model. Recall that some regions, such as 4, 5, 6, and 7, are not densely populated with humans across time compared to other regions (such as 1, 2, 3, and 10). The POPP and POPP-Beta models, however, still managed to improve the percentage of positive observations. This shows that the models correctly predicted that people would be present in particular locations at partic-




Fig. 19 This graph shows the percentage of time that the robot observed activities when it was present in a region. It is a measure of how successful the robot’s visit policy (choice of visit time and visit location) was in finding people. It presents results for for the FOPP, POPP and POPP-Beta algorithms

ular times. One should note that region 6 contains vending machines which are often detected as a person by the upper body detector. This leads to the FOPP model planning to visit this particular location when no activity is taking place. The POPP and the POPP-Beta models were able to correct the miscounts occurring in region 6, providing a better estimate of the posterior over the arrival rate λ . This leads to models that better capture the true underlying process and thus support more accurate exploration-exploitation trade-offs.

During the first few days of each 15 day phase the robot primarily explores since each model initially has a highly uncertain estimate of λ . Every three days, the sensor model for both the POPP and the POPP-Beta is updated to represent more accurate true positive and false positive rates. Note that the first 3 days of each exploration for the POPP and the POPP-Beta, the sensor model is set to follow a perfect sensor model (i.e. setting the true positive rate to 1.0 and false positive rate to 0.0) with the hyperparameters for the POPP-Beta set to follow uniform distribution. As more days of data are experienced the estimates increase in confidence and the robot starts to exploit this increased confidence by visiting locations which are likely to provide higher counts⁸. Figure 20 shows the number of actual humans (performing some activities) observed throughout each its exploration policy. Looking at the raw numbers provided by the figure, it seems that the exploration policy following the POPP model performed really well in finding people. However, this metric is quite unfair due to the time when the exploration was conducted (e.g. the exploration policy following the POPP-Beta model was conducted in Saturday and Sunday in some occasions). To allow us to produce a metric for a fair comparison across three models (FOPP, POPP, and POPP-Beta) deployed at different times (and thus experiencing different population

dynamics), we look at the ratio between the expected observations made by a baseline policy and those made by our exploration policy in the same period. To create the baseline total for each model we take the true counts experienced for its first three days then multiply these by five to give an expected total over 15 days (the number of days of data available to every model). Three days are chosen to align with the 3-day periodical update of the sensor model for the POPP and the POPP-Beta and create a uniform baseline across different explorations since both the POPP and the POPP-Beta explorations should act similarly to the FOPP exploration in the first 3 days. This is the denominator in Eqn. 19, where $s(n)$ is the (true) number of people observed on day n . This is used to divide the cumulative number of observations up to the current day:

$$\hat{s}(n) = \frac{\sum_{i=1}^n s(i)}{\sum_{i=1}^3 s(i) * 5} \quad (19)$$

Given this, a \hat{s} score of 1.0 on day 15 shows that people have been observed people at the rate of the baseline, i.e. the underlying model has failed to exploit additional data correctly. A result over 1.0 shows that the model has exploited the available data to observe people at a greater rate than in the first 3 days. Figure 21 presents the cumulative normalised true counts of people observed by the robot across the three phases. This shows that exploration driven by the POPP and the POPP-Beta models improves the number of people observed during these phases. By the end of each of these two phases, the ratio is around 1.7. On the other hand, the FOPP showed a stable ratio around the baseline (1.0 at day 15), this means that the FOPP is not be able to improve the number of people observed over time.

Exploration limitations

There are some limitations in our robot exploration experiment as a result of carrying them out on a real robot in an uncontrolled setting. A major limitation is how we stretched our assumption about (approximately) constant population dynamics throughout the 69-day exploration. It is clear that the population in a university building fluctuates within a single academic semester (e.g. students tend to be around in the middle of the semester—Fig. 20—than during the last week of the semester—Fig. 20). We mitigate this by normalising the raw counts following Eq. 19. The normalisation is effective when the differences in population dynamics are within reasonable range.

Due to the impact of occasional hardware and sensor malfunctions, plus an external time limit for running the

⁸ Note that this change from exploration to exploitation occurs naturally and gradually in an upper bound-based model, and therefore the characterisation of the behaviour as exploring or exploiting is a *post-hoc* justification.

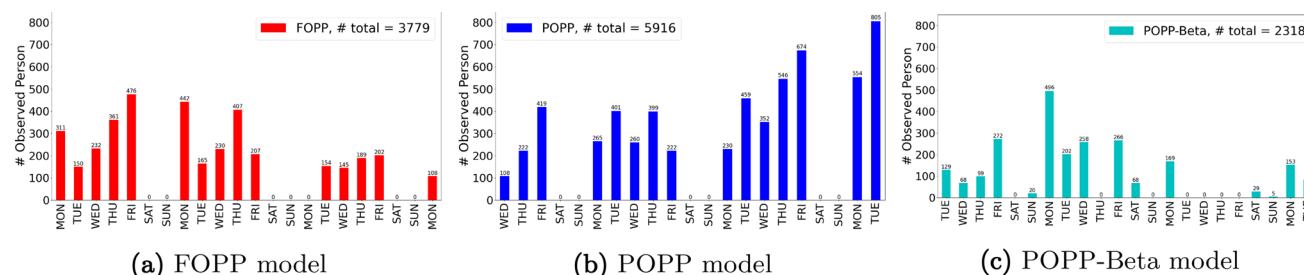


Fig. 20 The number of actual activities observed by the robot for each exploration policy




Fig. 21 The improvement ratio of activity observations during each phase of the trial. The dash line indicates a baseline performance, i.e., no improvement in exploration over time

experiments (winter break was approaching), we further stretched our assumption by assuming the population dynamics during weekends were similar to weekdays. However, upon inspection, there were far fewer students on the weekends for the robot to observe than on the weekdays (see Fig. 20). The mismatch between our assumption and the experimental setting mostly affected exploration following the POPP-Beta policy, which was tested last. The inclusion of weekends creates a big deviation in the population dynamics that renders the normalisation ineffective. Although the POPP-Beta policy was affected by a large variation in population dynamics, the policy was still able to improve the number of people observed over time. This can be shown by removing the weekends from calculations in Eq. 19, resulting 1.48 at day 15 for the POPP-Beta policy. An ideal experiment would be to run robot explorations with multiple identical robots employing different exploration policies at the same time.

9 Conclusion

This article has presented Bayesian estimators for (1) estimating human activities, as count data, at each time of the day and (2) helping an autonomous robot optimise between exploring for new time-place combinations where it might discover a high-level of human activity and re-visiting learnt time-place combinations that a wealth of human activity is

to be found. Our work was motivated by the application of counting people from an autonomous mobile robot using noisy sensors and perception algorithms. The work extends our prior work Jovan et al. (2018) with two main contributions. First, we presented variations of our previous POPP formulation: POPP-Beta extends POPP by accounting for the unreliability of the observation model; C-POPP extends POPP by modelling the case when sensors are uncorrelated; and POPP-Dirichlet combines POPP-Beta and C-POPP to provide the benefits of each correction. Evaluations on synthetic data and observations taken by a robot show that each extension provides progressively more accurate estimates than the POPP filter. Second, posteriors from FOPP, POPP and POPP-Beta were used to drive exploration by a mobile robot for a series of three exploration experiments. An upper bound interval exploration method in combination with Fourier transformation was used to solve the exploration-exploitation problem. This resulted in a labelled dataset of human presence counts. Our initial evaluation demonstrated that POPP and POPP-Beta were able to drive the robot to observe more people over time than the FOPP-based method.

There are many directions for further work including utilizing C-POPP and POPP-Dirichlet to drive the robot observation in an extended time period, allowing another filter strategies for faster and more accurate posterior estimates, and removing convenient closed forms of conjugate priors in the sensor models.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10514-022-10070-9>.

Funding The funding was provided by the FP7 Information and Communication Technologies; Engineering and Physical Sciences Research Council

Declarations

Conflict of interest The Author(s) declare(s) that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as

long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Sensor model

Here we provide detailed sensor models for each region. The sensor model was built from the first 15 days of data to give a representative idea how each sensor performs across regions

Table 3 Averaged sensor model for each region trained from 15 days of data

Region	Sensor	True Negative	True Positive
1	Leg	0.820	0.102
	Upper body	0.749	0.244
	Scenery change	0.760	0.612
2	Leg	0.991	0.655
	Upper body	0.862	0.691
	Scenery change	0.826	0.778
3	Leg	0.854	0.116
	Upper body	0.833	0.130
	Scenery change	0.780	0.687
4	Leg	0.896	0.180
	Upper body	0.967	0.227
	Scenery change	0.897	0.592
5	Leg	0.918	0.086
	Upper body	0.881	0.200
	Scenery change	0.877	0.957
6	Leg	0.964	0.351
	Upper body	0.929	0.143
	Scenery change	0.803	0.541
7	Leg	0.949	0.264
	Upper body	0.829	0.071
	Scenery change	0.939	0.090
8	Leg	0.889	0.473
	Upper body	0.791	0.360
	Scenery change	0.900	0.591
9	Leg	0.702	0.383
	Upper body	0.711	0.172
	Scenery change	0.591	0.673
10	Leg	0.956	0.537
	Upper body	0.973	0.423
	Scenery change	0.823	0.584

for each of our experiments. For the experiment in Sect. 7, we used 48 days of data, and as we performed four fold cross validation, we only used 12 days of data to build the sensor model and the rest (36 days) of data were used to test and build the POPP models and its variants. During the exploration setting in Sect. 8, only the first 3 days from 15 days of exploration for each exploration model were used to build the sensor model Table 3.

References

- Alba, E., & Dorronsoro, B. (2005). The exploration/exploitation trade-off in dynamic cellular genetic algorithms. *IEEE Transactions on Evolutionary Computation*, 9(2), 126–142.
- Audibert, J. Y., & Munos, R., & Szepesvári, C. (2009). Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19), 1876–1902.
- Bratcher, T. L., & Stamey, J. D. (2002). Estimation of poisson rates with misclassified counts. *Biometrical Journal*, 44(8), 946–956.
- Coppola, C., Krajník, T., Duckett, T., & Bellotto, N., et al. (2016). Learning temporal context for activity recognition. European conference on artificial intelligence (pp. 107–115).
- Dondrup, C., Bellotto, N., Jovan, F., & Hanheide, M., et al. (2015). Real-time multisensor people tracking for human-robot spatial interaction. Proceedings of the 2015 international conference on robotics and automation. ICRA/IEEE. Duckworth, P., Gatsoulis, Y., Jovan, F., Hawes,
- Hogg, N., & Cohn, D. A. (2016). Unsupervised learning of qualitative motion behaviours by a mobile robot. Unsupervised learning of qualitative motion behaviours by a mobile robot. (p. 1043–1051). Association for Computing Machinery.
- Dvorzak, M., & Wagner, H. (2016). Sparse bayesian modelling of under-reported count data. *Statistical Modelling*, 16 (1), 24–46.
- Hanheide, M., Göbelbecker, M., Horn, G.S., Pronobis, A., Sjöö, K., Aydemir, A., others (2017). Robot task planning and explanation in open and uncertain worlds. *Artificial Intelligence*, 247, 119–150.
- Hawes, N., Burbridge, C., Jovan, F., Kunze, L., Lacerda, B., & Mudrová, L., Young, J., Wyatt, J., Hebesberger, D., Kortner, T., Ambrus, R. (2016). The strands project: Long-term autonomy in everyday environments. arXiv preprint [arXiv:1604.04384](https://arxiv.org/abs/1604.04384).
- Jovan, F., Wyatt, J., & Hawes, N. (2018, 09–11 Apr). Efficient bayesian methods for counting processes in partially observable environments. A. Storkey & F. Perez-Cruz (Eds.), Proceedings of the twenty-first international conference on artificial intelligence and statistics (Vol. 84, pp. 1906–1913). Playa Blanca, Lanzarote, Canary Islands: PMLR.
- Jovan, F., Wyatt, J., Hawes, N., & Krajník, T. (2016). *A Poisson-Spectral Model for Modelling the Spatio-Temporal Patterns in Human Data Observed by a Robot*. Daejeon, South Korea: Pro 1000 conf on intelligent robots and systems.
- Kaplow, R., Atrash, A., & Pineau, J. (2010). Variable resolution decomposition for robotic navigation under a pomdp framework. 2010 IEEE international conference on robotics and automation (pp. 369–376).
- Martinez-Cantin, R., De Freitas, N., Brochu, E., Castellanos, J., Doucet, A. (2009). A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, 27 (2), 93–103.
- Molina, S., Cielniak, G., Duckett, T. (2022). Robotic exploration for learning human motion patterns. *IEEE Transactions on Robotics*, 38 (2), 1304–1318. DOI: 10.1109/TRO.2021.3101358.

- Papoulis, A., & Pillai, S.U. (2002). Probability, random variables, and stochastic processes. Tata McGraw-Hill Education.
- Pineau, J., Montemerlo, M., Pollack, M., Roy, N., Thrun, S. (2003). Towards robotic assistants in nursing homes: Challenges and results. *Robotics and autonomous systems*, 42 (3–4), 271–281.
- Spосто, R., Preston, D.L., Shimizu, Y., Mabuchi, K. (1992). The effect of diagnostic misclassification on non-cancer and cancer mortality dose response in a-bomb survivors. *Biometrics*, 48 (2), 605–617.
- Sridharan, M., Gelfond, M., Zhang, S., Wyatt, J. (2019). Reba: A refinement-based architecture for knowledge representation and reasoning in robotics. *Journal of Artificial Intelligence Research*, 65, 87–180.
- Stamey, J.D., & Young, D.M. (2005). Maximum likelihood estimation for a poisson rat parameter with misclassified counts. *Australian & New Zealand Journal of Statistics*, 47 (2), 163–172.
- Stamey, J.D., Young, D.M., Bratcher, T.L. (2004). Bayesian predictive probability functions for count data that are subject to misclassification. *Biometrical Journal*, 46 (5), 572–578.
- Street, C., Lacerda, B., Mühlig, M., & Hawes, N. (2020). Multi-robot planning under uncertainty with congestion-aware models. 19th international conference on autonomous agents and multiagent systems (aamas). Auckland, New Zealand.
- Tenenbein, A. (1970). A double sampling scheme for estimating from binomial data with misclassifications. *Journal of the American Statistical Association*, 65 (331), 1350–1361.
- Velez, J., Hemann, G., Huang, A.S., Posner, I., Roy, N. (2012). Modelling observation correlations for active exploration and robust object detection. *Journal of Artificial Intelligence Research*, 44, 423–453.
- Whittemore, A.S., & Gong, G. (1991). Poisson regression with misclassified counts: Application to cervical cancer mortality rates. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 40(1), 81–93.
- Winkelmann, R., & Zimmermann, K. F. (1993). *Poisson-logistic regression*. München: Volkswirtschaftl. Fakultät d. Ludwig-Maximilians-Univ.
- Wyatt, J. (1998). Exploration and inference in learning from reinforcement

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Ferdian Jovan is a research scientist at University of Bristol. He completed his Ph.D from University of Birmingham in 2018 on modelling temporal patterns of human occupancy from robots' unreliable sensors and then worked at the Department of Engineering Science, University of Oxford to develop generative models for predicting battery health and its remaining useful life on real-world data. In 2019, he joined the artificial intelligent lab led by Prof. Sara Bernardini

where he developed multi-agent planning system for inspection, maintenance and repair of wind turbines using multi-platform autonomous robots. His current research focuses on machine learning approaches to intelligent systems including mobile robotics and autonomous agents.



PhD research focuses in planning, decision making and inferences under uncertainty.



and autonomous robots. With this focus, he has secured and led many UK and European projects including STRANDS, ALOOF and ADE. He currently leads the GOALS research group that focuses on long-term autonomy and task and mission planning for mobile robots which must operate for extended periods (days, weeks or months) in dynamic, uncertain environments.



GeRT on robot manipulation, and CoSy on cognitive robotics. His current work, in Amazon, focuses on robot manipulation and mobile fleet management for increasing the efficiency of mobile fleet in production.

Milan Tomy is a Ph.D candidate from University of Sydney. She graduated from Electronics and Communication Engineering, National Institute of Technology Calicut in 2016 and received her Master degree in Robotics from University of Birmingham in 2017. After three years of being a research scientist in several industrial companies, she continued pursuing her PhD degree. She is most fascinated by the technology that enables machines to understand and aid humans better. Her

Nick Hawes is an associate professor of robotics at the Oxford Robotics Institute, University of Oxford. He received his Ph.D from University of Birmingham in 2003, and then worked at the Media Lab Europe as a Research Fellow. In 2004, he came back to University of Birmingham as a Research Fellow, became a lecturer in 2009, and a reader in Autonomous Intelligent Robotics in 2015. His research focuses on Artificial Intelligence (AI) for long-term autonomy in intelligent

Jeremy Wyatt is an honorary professor of Robotics and Artificial Intelligence, University of Birmingham, and the General Manager of Robotics AI, Amazon. He did his Ph.D in Artificial Intelligence from University of Edinburgh. He then joined University of Birmingham as a Lecturer and obtained a full professorship in 2013. In University of Birmingham, he led and coordinated many UK and European projects including PacMan on robot manipulation, CogX on robot planning,