Cox Bayesian Optimization for Police Patrolling

Roman Marchant, Di Lu and Sally Cripps
Centre for Translational Data Science
The University of Sydney
Australia NSW 2006
{roman.marchant,di.lu,sally.cripps}@sydney.edu.au

Abstract

This paper presents a generalization of *Bayesian Optimization* (BO) that can cope with objective functions defined over discrete spaces. Typically, BO assumes a latent model represented by a *Gaussian Process* (GP) prior under Gaussian likelihood assumptions. This works well for continuous real valued objective functions that are Lipschitz continuous. However, discontinuous objective functions, which are a consequence of discrete inputs and outputs, present important challenges. The contribution in this paper is to address these challenges by making use of a Log Gaussian Cox Process (LGCP) in the BO framework, which has a GP as an argument of the link function, and uses *Monte Carlo Tree Search* (MCTS) to deal with discrete inputs. We evaluate the proposed algorithm over a synthetic objective function and apply it to a real world problem – police patrolling, where observations are integers at a set of discrete locations.

1 Introduction

Over the last decade, *Bayesian Optimization* (BO) [3, 14] has gained popularity among the Machine Learning community due to its flexibility and efficiency for solving expensive optimization problems, spanning a wide range of areas in academia and industry. Some interesting applications of BO include environmental monitoring [9], experimental design [2], and interactive user interface [3]. In this paper, we propose *Cox BO* (CBO), which allows BO to optimise functions with observations arising from the realisation of discrete events. In the spatial-temporal scenario, this can be represented by an in-homogeneous Poisson Point Process, whose intensity is commonly modelled by a non parametric Bayesian model, particularly a GP, giving rise to the popular Log Gaussian Cox Process [11, 6].

Various methods have been proposed to approximate posterior distributions with non-Gaussian likelihoods, such as Laplace Approximation [13], Expectation Propagation [13] and *Variational Bayes* (VB) [12]. A unifying framework for constructing flexible models using a generic exponential family likelihood was proposed by [15] and defined as *Generalized Gaussian Process Model* (GGPM). GGPM assumes y follows an exponential distribution where the expected value of y, denoted by $\eta = g^{-1}(f(\mathbf{x}))$, where g is a link function, varies according to the distribution of y. If y is Poisson then $g(\mathbf{x}) = log(\mathbf{x})$. The likelihood function in LGCPs lacks of closed form solution, and there has been seen significant research into the efficient estimation of the intensity function given possibly large number of observations, using VB [8, 1] and other Bayesian inference techniques [16]. Despite latest developments in BO and the existing LGCP theory, to the best of our knowledge, there does not exist a piece of work which bridges both areas. Our paper describes a framework for Bayesian Optimization to handle a Poisson likelihood and build on the work by [8, 15, 14].

An important contribution of this paper is the application of the proposed methodology to monitoring spatial temporal event data. Particularly, we have identified the problem of optimal police patrolling to better understand and uncover areas with higher number of crimes. In the more general setting,

Workshop on Modeling and Decision-Making in the Spatiotemporal Domain, 32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada.

one would ideally build a continuous representation of λ that can accommodate multiple patrolling modalities (helicopters, foot, drone, car, etc). However, in this paper, we have discretised the action space over street networks to reflect car patrolling. We deal with this discrete input space by applying a *Monte Carlo Tree Search* (MCTS) [4] approximation in the Sequential BO setting described by [10].

Presently, there is an ongoing concern about over-policing resulting from maximising only the expected number of crimes. By using the following technique, it is possible to reduce the bias in the decision making process by taking into account the uncertainty over these estimations.

2 Cox Bayesian Optimization

Cox Bayesian Optimization (CoxBO) is an adaptation of vanilla BO that is able to deal with discrete event data over a continuous space. For a detailed background on BO we refer the reader to the comprehensive review by [14]. In this paper we discretise the input space and aggregate observations into counts of events, y, which breaks the Gaussian assumption in vanilla BO. A statistical model for the data generating process is given by

$$y_n(\mathbf{x}) \sim \text{Poisson}(\lambda(\mathbf{x}_n))) \ , \ \eta(\mathbf{x}) = \log \lambda(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}')) \ . \tag{1}$$

where $\lambda(\mathbf{x}_n)$ represents the input-dependent intensity for the in-homogeneous point process used to describe y_n. Following the principles behind a LGCP model, CoxBO places a latent model over the log of the intensity function. Since the goal is to find the locations for which the counts y are higher, the log-intensity will be used as a proxy optimization routine.

A critical characteristic of every BO algorithms is the acquisition function $u(\mathbf{x}) \in \mathbb{R}$, which is responsible for the active learning nature of the solution and guides the search for the location of the optimum. As in vanilla BO, the acquisition function $u(\mathbf{x})$ is designed to be rapidly evaluated though non-convex auxiliary goal function, effectively minimizing the number of samples and automatically controlling the balance between exploitation and exploration. Algorithm 1 presents the CoxBO algorithm, which incrementally gathers samples from the goal function and uses active decision making to select the locations of where to evaluate next. The data-set $\mathcal{D}$ represents the set of N location-sample pairs, by $\mathcal{D} = \{\mathbf{x}_n, y_n\}_{i=1...N}$.

Algorithm 1 Cox Bayesian Optimization (CoxBO)

procedure $BO(f, u, \text{GP}, T)$
  for $t = 1, 2, ..., T$ **do**
    Find $\mathbf{x}_t = \text{argmax}_{\mathbf{x}}\, u(\mathbf{x}|D_{1:t-1})$
    Obtain a realization of the objective function: $y_t \sim p(y|\log \lambda(\mathbf{x}_t))$
    Augment the observed data-set with new data: $D_{1:t} = D_{1:t-1} \cup \{(x_t, y_t)\}$
    Update the GP model of $\eta(\mathbf{x})$

In practice, the hyper-parameters of the GP model are pre-trained with existing data and can be re-trained over time depending on available computational resources. Since the likelihood is not tractable and there are a large number of observations, we have conducted inference by using Stochastic Variational Inference [7, 1, 8]. The predictive distribution for y at $\mathbf{x}^\star$, $p(y^\star|\mathbf{x})$ is computed in two stages: first predicting the latent function by marginalizing over all possible function values,

$$p(\eta^\star|\mathbf{x}^\star, \mathcal{D}) = \int p(\eta^\star|\boldsymbol{\eta}, \mathbf{X}, \mathbf{x}^\star)p(\boldsymbol{\eta}|\mathcal{D})d\boldsymbol{\eta} \ , \tag{2}$$

then estimating the actual prediction over y by using the Poisson likelihood function, which becomes

$$p(y^\star|\mathbf{x}^\star, \mathcal{D}) = \int p(y^\star|\eta^\star)p(\eta^\star|\mathbf{x}^\star, \mathcal{D})d\eta^\star \ . \tag{3}$$

The acquisition function is defined in terms of the predictive distribution for the latent function, given by Equation 2. Since the log-intensity is modelled with a GP, as in Equation 1, it is possible to evaluate the posterior mean estimate $\mu(\mathbf{x}^\star)$ and the variance $\sigma^2(\mathbf{x}^\star)$ for $\log \lambda(\mathbf{x}^\star)$. Similar to other BO variantes, here we allow for any acquisition function $u(\mathbf{x})$ [Shahriari2015].

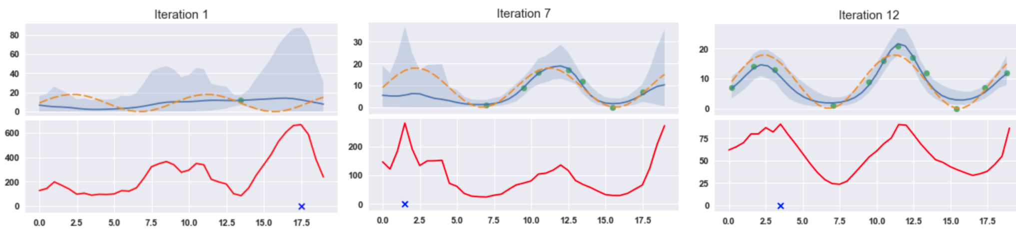

Figure 1: Illustration of CoxBO over 1D data. Yellow dashed curve (ground truth) – green dots (observed counts) – blue line (mean estimate for intensity) – shaded area (95% confidence intervals) – red line (acquisition function) – blue cross (maximum of the acquisition function).

In the more generic setting, observations correspond to events in a space-time continuum (or any other domain). For the first iteration of this work we discretise the space $\mathbf{x}$, where y_i represents a count of events over cell i in the space.

We avoid the use of myopic approximations and implement sequential setting proposed by [10]. Effectively, this leads to evaluating the objective function at a fixed set of locations. A *Partially Observable Markov Decision Process* (POMDP) is defined as an analog to the discrete-action BO scenario, which is solved using MCTS.

3 Experiments

3.1 Synthetic Poisson Process Data

The first experiment illustrates the behaviour of CoxBO for synthetic data with known ground truth. For evaluation purposes, we use a known log-intensity function $\lambda(x) = 9\sin(0.7x) + 9$. The input space for x corresponds to a one-dimensional space $x \in [0, 19]$. We executed Algorithm 1 using an UCB [5] acquisition function, with parameter $\kappa = 7.5$. As with most BO techniques, this parameter was manually tuned and represents the trade off between exploration-exploitation.

The results, illustrated in Figure 1, show that the application of CoxBO to this data is highly beneficial in multiple ways. Firstly, it can be appreciated that the model has converged after 9 iterations, which is highly beneficial for the evaluation of costly functions. It is also worth mentioning that the uncertainty over the counts remains positive for any x, given the real valued exponential link function. I.e. even if the estimate of the GP provides a negative estimate, the final sample for the counts y is a positive integer.

3.2 Crime Simulation over Street Networks

The ultimate goal is to better understand the areas for which a higher number of crimes occur, while exploring the latent function of crime intensity across space. The active search procedure corresponds to decision making over street networks, where the counts (aggregated observations) are at the cells and actions are along an interconnected graph of streets.

The objective function is the number of crimes (Domestic Violence related assaults) in a given area for the city of Sydney in New South Wales, Australia. The data are 467 crime occurrences, within a one month period – December 2014. Figure 2 shows a ground truth density from the raw crime locations. The environment was discretised into 425 cells and the street network graph was obtained by using Open Street Maps API.

The active learning nature of the CoxBO algorithm over count data was used to decide actions of a police patrolling car within the street network. Actions correspond to driving along edges of the graph, highlighting the discrete notion of the input domain, and observations y corresponds to a number of crimes observed in the neighbourhood region. The covariance function used to represent the latent function for the log-intentisy of the GP was the Matern32, for with hyper-parameters were estimated from random samples of the environment and result in $\ell = 0.156$ and $\sigma_f = 7.732$. Similar to 3.1 we use an UCB acquisition function $u(\mathbf{x})$. The goal is to identify the changing behavioural patterns of the patrolling, in terms of exploration and exploitation.

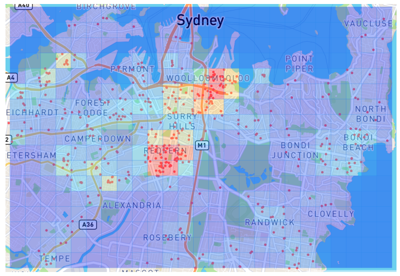

Figure 2: Fit Poisson distributed data using Gaussian Process regression.

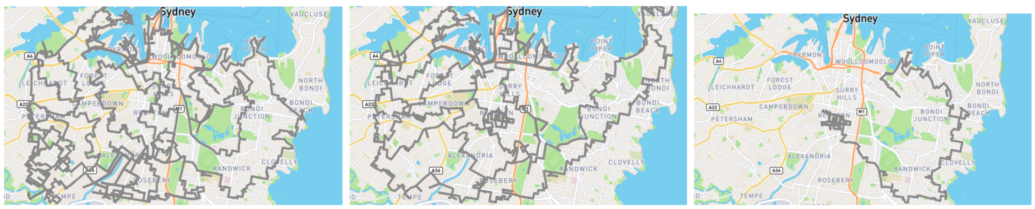

Figure 3: Resulting path for κ equal to 50, 8.5 and 2 respectively.

At each time-step, the maximization of the acquisition function was conducted by building an approximated solution using MCTS. The parameters for MCTS are a 4 step look-ahead, with $\beta = 50$, where β is the exploration exploitation trade-off in the expansion function UCT for growing the tree. We ran 400 simulations per experiment.

It can be seen that the resulting paths are considerably different between each other and serve different purposes. Figure 3 illustrates the effect of modifying the exploration-exploitation trade off parameter from the acquisition function, κ. A value of $\kappa = 50$ (left in Figure 3), results in a pure exploration behaviour, i.e. a maximum coverage scenario, where not much importance is given to the expected value of crime. For a value of $\kappa = 2$ there is an over-importance to the expected value of crime intensity and the patrolling car gets stuck and known high risk locations without exploring. $\kappa = 8.5$, shown in the centre of Figure 3, presents a more sensible and parsimonious approach for simultaneously dealing with exploration and exploitation. It can be seen how the police patrolling car meets two of the most important requirements: exploitation around areas with expected high crime and also exploration to achieve discovery of other height crime areas.

By taking an evidence based approach for decision making under uncertainty, and particularly finding a convenient value for κ, it is possible to reduce the existing bias and over-policing. The reinforcing phenomena between over-policing and data bias can get reduced and potentially removed if the trade off between exploration and exploitation is handled in a principled manner.

4 Conclusion

This paper presents Cox Bayesian Optimization, which corresponds to the generalization of BO for the case when the objective function is defined over the positive integers, following a Poisson distribution. We use Stochastic Variational Inference, which coupled with a logarithmic link function and MCTS give rise to practical implementation of CoxBO that is demonstrated to work on real data-set. We have validated the methodology on synthetic data and on a real scenario consisting of a police patrolling environment.

We envisage several areas for future work that could increase the range of applications of this work. Particularly expand the models to generic exponential distributions with multidimensional parameters paces by using muti-output GsP. Another area of improvement is to model the data generation process following a mixture of Poisson distributions, which will leverage the strong limitation of identical mean and variance arisigin from a single Poisson distribution.

References

[1] R. P. Adams, I. Murray, and D. J. C. MacKay. Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities. *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, pages 1–8, 2009. ISSN 15324435. doi: 10.1145/1553374.1553376. URL `http://portal.acm.org/citation.cfm?doid=1553374.1553376`.

[2] J. Azimi, A. Jalali, and X. Fern. Hybrid batch bayesian optimization. *CoRR*, abs/1202.5597, 2012. URL `http://arxiv.org/abs/1202.5597`.

[3] E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report arXiv:1012.2599, University of British Columbia, 2010.

[4] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 2012.

[5] D. Cox and S. John. A Statistical Method for Global Optimization. In *IEEE Conference on Systems, Man and Cybernetics*, 1992.

[6] P. Diggle, P. Moraga, B. Rowlingson, and B. Taylor. Spatial and spatio-temporal log-gaussian cox processes: Extending the geostatistical paradigm. *Statistical Science*, 28(4):542–563, 11 2013. ISSN 0883-4237. doi: 10.1214/13-STS441.

[7] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic Variational Inference. *Journal of Machine Learning Research*, 2013.

[8] C. Lloyd, T. Gunter, M. A. Osborne, and S. J. Roberts. Variational Inference for Gaussian Process Modulated Poisson Processes. In *International Conference on Machine Learning (ICML)*, 2015.

[9] R. Marchant and F. Ramos. Bayesian Optimisation for Intelligent Environmental Monitoring. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2012.

[10] R. Marchant, F. Ramos, and S. Sanner. Sequential Bayesian Optimisation for Spatial-Temporal Monitoring. In *Uncertainty in Artificial Intelligence (UAI)*, 2014.

[11] J. Møller, A. R. Syversveen, and R. P. Waagepetersen. Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25(3):451–482, 1998. ISSN 03036898. doi: 10.1111/1467-9469.00115.

[12] M. Opper and C. Archambeau. The Variational Gaussian Approximation Revisited. *Neural Computation*, 21(3):786–792, 2009. ISSN 0899-7667. doi: 10.1162/neco.2008.08-07-592. URL `http://www.mitpressjournals.org/doi/10.1162/neco.2008.08-07-592`.

[13] C. E. Rasmussen and C. Williams. *Gaussian processes for machine learning*. The MIT Press, Cambridge, Massachusetts, 2006.

[14] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. D. Freitas. Taking the Human Out of the Loop:A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):1–24, 2015. ISSN 1098-6596. doi: 10.1017/CBO9781107415324.004. URL `http://www.cs.ox.ac.uk/people/nando.defreitas/publications/BayesOptLoop.pdf`.

[15] L. Shang and A. B. Chan. On Approximate Inference for Generalized Gaussian Process Models. 2013.

[16] M. Teng, F. S. Nathoo, and T. D. Johnson. Bayesian Computation for Log-Gaussian Cox Processes – A Comparative Analytis of Methods. *Journal of Statistical Computation and Simulation*, 2017.

5 Acknowledgment

This work was supported by the *Sydney Informatics Hub* (SIH) at The University of Sydney and the data was provided by the New South Wales Police Force.