Adaptive Sensor Placement for Continuous Spaces

James A Grant, Alexis Boukouvalas, Ryan-Rhys Griffiths,
David S Leslie, Sattar Vakili, Enrique Munoz de Cote

Lancaster University

PROBLEM

- Placing sensors to detect events of interest,
  - **Maximise number of events detected** minus cost.
- Events arise according to a **Non-homogeneous Poisson process**.
- We are interested in a sequential version of the problem,
  - **Continuum-Armed Bandit**.

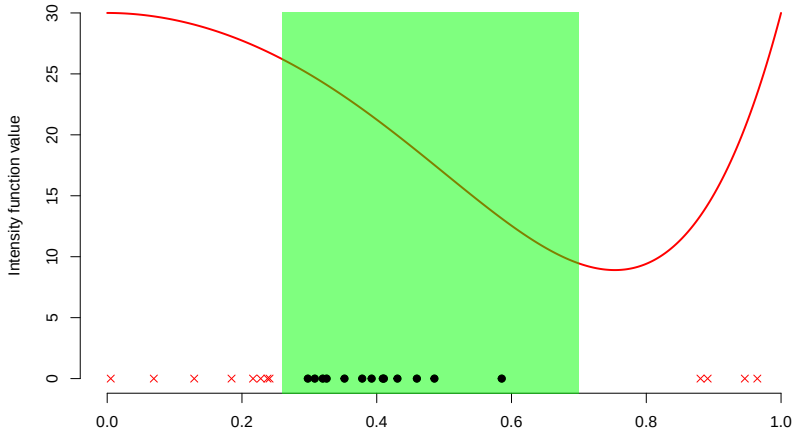

Events generated according to a Poisson process. The green interval is the selected sensing region.

CHALLENGES

We consider the **regret minimisation framework** and require an approach which tackles the following challenges:

- **Scalable inference**
  - Dependence on number of observed events important
- **Continuous action space**
  - Determining the best amongst infinitely many actions
- **Appropriate exploration/exploitation**
  - UCB, TS etc. - need to be adapted to point process data.

SOLUTION

We propose an approach which meets the challenges by use of

- **Bayesian histogram**
  - Efficient nonparametric estimation, asymptotically optimal error shrinkage
- **Progressive discretisation**
  - Both of histogram and action space.
- **Thompson Sampling**
  - Readily deployable without tuning UCBs

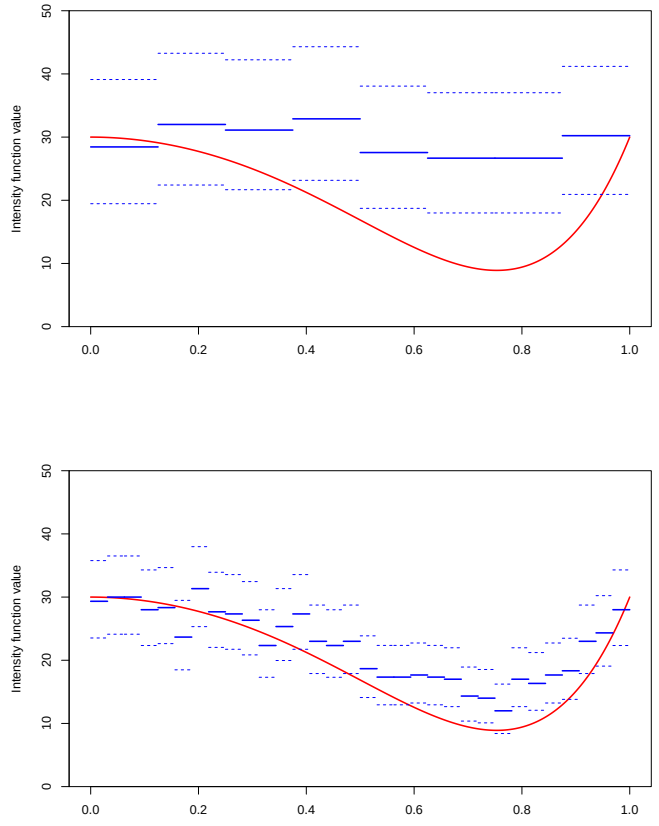

Bayesian histogram confidence intervals, showing progressive discretisation

Lancaster University

RESULTS

We have a **bound on the Bayesian Regret** of order $\tilde{O}(T^{2/3})$, and **strong empirical performance**:

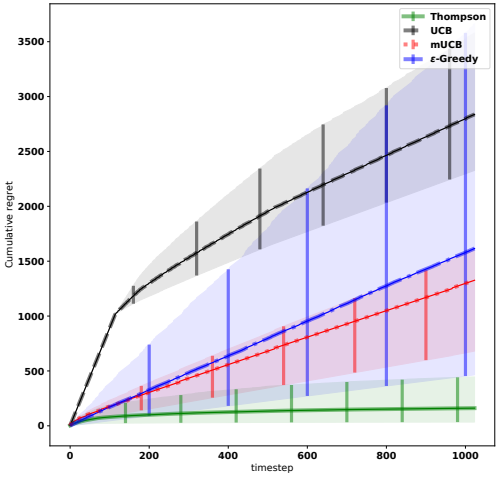

Regret accumulated by Thompson Sampling, and competitors

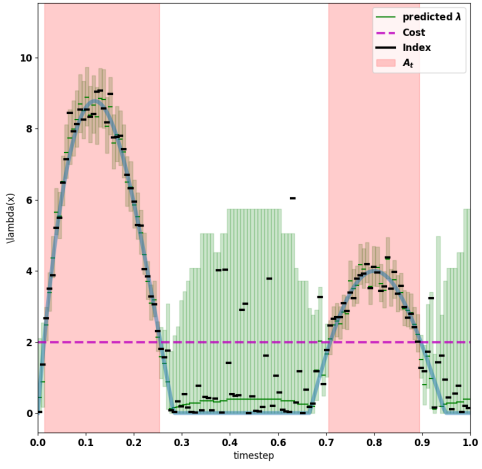

Posterior distribution, round 900 with Thompson Sampling

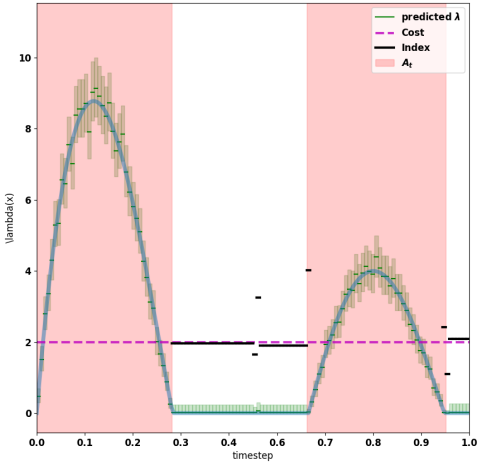

Posterior distribution, round 900 with UCB approach

Lancaster University

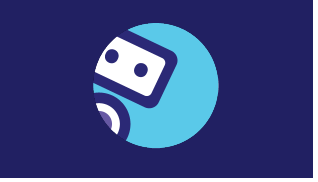

#167