

# Bayesian Optimization for real-time $\gamma$ -localization measurements with robotic platforms

Ioannis Tsitsimpelis<sup>1\*</sup>, Andrew West<sup>2</sup>, Kartikey Mathur<sup>1</sup>, Anže Jazbec<sup>3</sup>, Luka Snoj<sup>3</sup>, Shengshu Liu<sup>2</sup>, Andrew Kennedy<sup>1</sup>, Francis R. Livens<sup>4</sup>, Barry Lennox<sup>2</sup>, C. James Taylor<sup>1</sup>, and Malcolm J. Joyce<sup>1</sup>

<sup>1</sup> School of Engineering, Lancaster University, United Kingdom

<sup>2</sup> School of Electrical and Electronic Engineering, University of Manchester, United Kingdom

<sup>3</sup> Research Centre for Radwaste Disposal, Williamson Research Centre and Dalton Nuclear Institute, University of Manchester, United Kingdom

<sup>4</sup> Reactor Physics Department, Jožef Stefan Institute, Slovenia

(\*) [i.tsitsimpelis3@lancaster.ac.uk](mailto:i.tsitsimpelis3@lancaster.ac.uk)

**Abstract**— An adaptive approach driven by Bayesian Optimization is described for applications where remote radiation measurements made with robots are constrained by stringent upper thresholds on the mass and power payload of the necessary instrumentation, as well as by the time window within which measurements must be made, ultimately affecting their quality and maximum area coverage. The algorithm presented in this paper is applied to a gimbal assembly that comprises a collimated cerium bromide detector to perform  $\gamma$ -localization. A Gaussian Process models the angular distribution of radiation, and measurement locations are dynamically selected via a composite acquisition function, which combines the Expected Improvement and Uncertainty Confidence Bounds functions. Convergence is driven by monitoring the rate of change of predictions and associated mean uncertainty. This approach enables accurate characterization of radiation fields while requiring up to 85% fewer measurements than conventional raster-type scanning. Its performance is evaluated in simulation, using previously obtained datasets used as measurement look-up tables, and validated in turn with hardware implementation, real-time scans.

**Keywords** — Remote measurements, Gamma localization, Bayesian Optimisation, Robotics.

## I. INTRODUCTION

Remote radiation measurements for monitoring and characterization are highly advantageous, as they mitigate the risks associated with manual assessments performed by radiation workers [1]. Additionally, these measurements can be conducted and augmented systematically, thereby enhancing the fidelity and confidence in the acquired knowledge of the target environment [2], [3]. However, their implementation can be particularly challenging due to stringent constraints related to accessibility and safety [4], [5]. These challenges are further constrained in applications where remote radiation measurements are conducted using robotic platforms, which are often limited by strict upper thresholds on the mass and power

payload of the necessary instrumentation; and operational constraints such as the time window available for data acquisition due to e.g. the finite battery life of the robot [6]. Within these limitations, selective sampling strategies can be particularly advantageous, as they enable the prioritization of measurement locations that maximize area coverage while minimizing resource expenditure within the aforementioned constraints.

This research applies a well-established Bayesian Optimization (BO) technique in conjunction with a  $\gamma$ -radiation localization apparatus specifically designed for deployment on mobile robotic platforms. The apparatus is a gimbal-collimated detector assembly that can be mounted atop a robotic platform to enable the estimation of radiation distributions within its surroundings. The system identifies the directional distribution of  $\gamma$ -ray counts, providing the robot with situational awareness of radiation in all directions. To overcome the long dwell times associated with raster-scans (where the detector is rotated sequentially at discrete intervals), a dynamic sampling approach is implemented, whereby Gaussian Process Regression (GPR) together with a weighted exploration-exploitation acquisition function are used iteratively to efficiently sample and predict the radiation distribution of the angular input space. The evaluations presented hereinafter were carried out both in simulation and experimentally. The results suggest that reliable distributions can be derived at a fraction of the measurements that would be normally required by a raster-type scans; requiring up to 85% fewer measurements than conventional raster-type scans.

The rest of the paper is outlined as follows: Section II describes the instrumentation and BO algorithm. Section III presents results from simulated and real-time implementations. Section IV discusses advantages, limitations, and further work plans, and section V presents the concluding remarks.

## II. METHODS

### A. Instrumentation

The radiation detection hardware comprises a miniature

cerium bromide (CeBr<sub>3</sub>) detector with a 10 x 10 mm crystal coupled to a photomultiplier tube and a miniature High Voltage module; hermetically enclosed in a cylindrical aluminium body (110 mm long, 25 mm diameter, 60 g, Figure 1a). The detector's body is coupled to a novel lightweight collimator [7], made from a series of 3D printed porous rings coated with tungsten powder (Figure 1b), while its face is covered by two 3D printed tungsten filament slot discs (Figure 1c). The assembly is mounted to a gimbal, designed and 3D printed in-house (Figure 1d), to enable angular scans in azimuth and incline. The signal processing is carried out by a RedPitaya Field Programmable Gate Array (FPGA, model: STEMLab 125-14) board configured as a Multi-channel Analyzer (MCA) (Figure 1e). The instrumentation processes are integrated with Python (version 3.13) and the ubiquitous Robot Operating System (ROS2) middleware. Figure 1f shows a picture of the gimbal on top of the payload area of a Clearpath robotics Jackal platform.

For every angular position, a separate energy spectrum is recorded, thus allowing the derivation of net and energy resolved angular responses (accumulated counts per angle) to a radiation field. The resulting distribution depends on the collimation shape and material, as well as the activity and geometrical characteristics of the radiation field (e.g. dispersal); and the scan field of the collimated detector. For instance, the horizontal angular response of a vertical slot collimator to the x-ray and photopeak associated with a <sup>137</sup>Cs point source located at the centre of the scan resembles the shapes depicted in Figure 1g. In general, the shape depicted for the x-ray (in red circles) is similar for up to ~200 keV, whilst above that threshold it appears as per the angular distribution of the photopeak depicted in black circles.

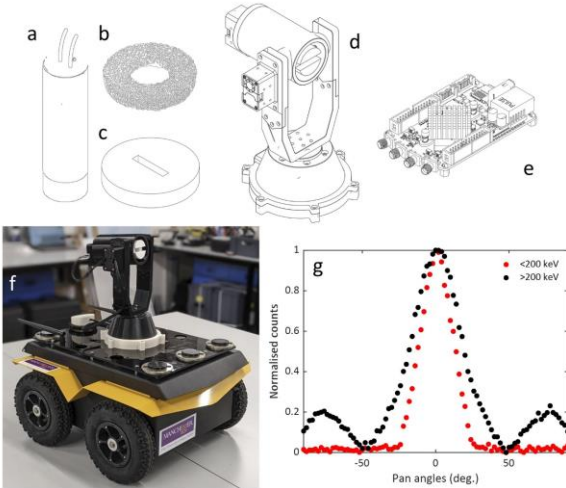


Figure 1: a) Drawing of the CeBr<sub>3</sub> detector used in this research; b) Drawing of a porous 3D printed plastic ring coated with tungsten powder. Seven of these are used to cover the body of the detector; c) Drawing of a 3D-printed tungsten slot collimator. Two of these are placed in front of the detector's face; d) Drawing of the custom-made gimbal that houses the detector; e) Drawing of the Red Pitaya FPGA board used to record spectral data; f) A picture of the gimbal (housing a mock-up collimator) mounted on the payload bay of the Jackal robotic platform; g) A plot showing the typical horizontal angular responses observed at low and high gamma energies when using a vertical slot lead collimator. Here, the responses are based on a simple geometric scenario where a point source is located at the centre of the scan.

## B. Bayesian Optimisation algorithm

The unknown radiation count distribution is modelled as a Gaussian Process (GP), with prior defined as:

$$f(\theta) \sim GP(\mu(\theta), k(\theta, \theta')) \quad 1$$

where  $\mu(\theta) = \mathbb{E}[f(\theta)]$  is the mean function (estimated radiation distribution) against  $\theta$  angles, and  $k(\theta, \theta')$  is the covariance function defined as a 3/2 Matérn kernel. This is defined as:

$$k(\theta, \theta') = \sigma_f^2 \left( 1 + \frac{\sqrt{3} |\theta - \theta'|}{l} \right) \exp \left( -\frac{\sqrt{3} |\theta - \theta'|}{l} \right) \quad 2$$

where  $\nu = 1.5$  is the smoothness control parameter,  $l$  is the length scale (in this case its value determines how quickly correlation decays at different angle distances), and  $\sigma_f^2$  is an amplitude scaling factor.

The algorithm initializes by requesting a set of measurements at the ends of the input space and  $n$  equally spaced angular locations within the input space, specified by the user. The initial dataset thus contains  $n$  observations, consisting of angular positions  $\theta_i$  and associated count measurements  $y_i$ . On the basis of this initial dataset, the covariance matrix  $\mathbf{K}$  is calculated (with entries  $\mathbf{K}_{i,j} = k(\theta_i, \theta_j)$ ) and subsequently used for the estimation of the radiation distribution across the whole input space. In particular, for each angular location with no recorded counts  $\theta^*$ , the GP posterior has mean and variance computed as:

$$\mu(\theta^*) = \mathbf{k}(\theta^*, \boldsymbol{\theta}) [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y} \quad 3$$

$$\sigma^2(\theta^*) = k(\theta^*, \theta^*) - \mathbf{k}(\theta^*, \boldsymbol{\theta}) [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{k}(\boldsymbol{\theta}, \theta^*) \quad 4$$

Where  $k(\theta^*, \boldsymbol{\theta})$ , being the cross-covariance function between the angle for which there is no associated count information, against all the angles for which this information exists;  $\sigma_n^2$  is a noise variance hyper-parameter;  $\mathbf{I}$  is the identity matrix;  $\mathbf{y}$  is a vector of the current measurements associated with  $\boldsymbol{\theta}$ ; and  $k(\theta^*, \theta^*)$  is the auto-covariance function.

Once (3) and (4) are computed, they are in turn utilized by a composite acquisition function  $\mathbf{w}(\boldsymbol{\theta})$ , comprising a slightly modified Expected Improvement (EI) function and the Upper Confidence Bounds (UCB), and defined as:

$$\mathbf{w}(\boldsymbol{\theta}) = \alpha \mathbf{EI}(\boldsymbol{\theta}) + (1 - \alpha) \mathbf{UCB}(\boldsymbol{\theta}) \quad 5$$

Where,  $\alpha$  is a weight parameter that determines the contribution of each component. The modified EI function is defined as:

$$\mathbf{EI}(\boldsymbol{\theta}) = (\boldsymbol{\mu}(\boldsymbol{\theta}) - f_{max} - \xi) \Phi(Z) + \boldsymbol{\sigma}(\boldsymbol{\theta}) \phi(Z) + \beta \boldsymbol{\sigma}(\boldsymbol{\theta}) \quad 6$$

Where:  $\xi$  is a small constant to encourage exploration;  $f_{max}$

is the maximum count value for the current locations for which associated counts exist;  $\sigma(\theta)$  is the standard deviation;  $\beta$  is the modification that scales the contribution of the standard deviation to bias more exploration,  $\Phi(Z)$  and  $\phi(Z)$  are the cumulative distribution function (CDF) and probability density function (PDF) of the standard normal distribution, respectively; and  $Z$  is defined as:

$$Z = \frac{\mu(\theta) - f_{max} - \xi}{\sigma(\theta) + \epsilon} \quad 7$$

With  $\epsilon$  a constant set to 1e-8 to avoid division by zero. The UCB function is designed to exploit areas of high activity and defined as:

$$UCB(\theta) = \mu(\theta) + \beta\sigma(\theta) \quad 8$$

Upon computation of  $w(\theta)$ , the next sampling location ( $\theta_{next}$ ) is selected on the basis of

$$\theta_{next} = \operatorname{argmax}(w(\theta)) \quad 9$$

In turn, a new measurement is carried out at the selected location and the model is updated. After a user-defined minimum number of iterations, convergence is evaluated on the basis of stability and uncertainty. For the former, the root mean square error is utilized, comparing the relative change of the current and previous predicted angular responses. For the latter, the mean estimated uncertainty is computed. These are defined as:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N \frac{(\mu(\theta_i) - \mu_{t-1}(\theta_i))^2}{\operatorname{Var}(\mu_{t-1}(\theta)) + \epsilon}} \quad 10$$

and

$$\frac{1}{N} \sum_{i=1}^N \sigma(\theta_i) \quad 11$$

When both metrics fall below user-defined thresholds, the algorithm terminates. The algorithm is implemented in Python using the Scikit-learn library.

### III. RESULTS

#### A. Simulation

In this simulation framework, previously obtained datasets function as look-up tables: every time the algorithm selects a measurement location, the corresponding recorded counts are retrieved from the dataset.

The black dots in Figure 2a collectively represent the gross spectral angular response of a vertically slotted, lead-collimated detector to a  $^{152}\text{Eu}$  point source positioned at the centre of the scan. The raster scan spans  $180^\circ$  in the horizontal plane, with  $2^\circ$  angular intervals, totalling 91 measurements. The orange diamond markers represent the initial measurement locations

requested during the initialization phase. The green triangles represent the measurements at locations subsequently selected by the algorithm. The dashed blue curve shows the final estimated response (predicted mean) across the input space, while the blue shaded area represents the associated uncertainty. In total, 14 measurement locations were "sampled" to approximate a radiation distribution that matches the raster scan ground truth very closely, requiring approximately 85% less measurements (14) than those carried out in the original raster scan (91).

Figure 2b presents a heatmap illustrating the progressive reduction of uncertainty along the input space as iterations increase. Similarly, Figure 2c displays a heatmap of the estimated radiation distribution at each iteration, with the top row showing the final estimate, also depicted in Figure 2a. Figure 2d visualizes the angle selection sequence through the combined acquisition heatmap: the bottom start of each blue vertical line represents the angle selected at each iteration, corresponding to the maximum value of the combined acquisition vector. Once selected, each angle is removed from the candidate set. It can be observed that the algorithm identifies and selects the central peak location early (prior to the second iteration of Gaussian Process Regression), and then alternates between sampling central and peripheral regions.

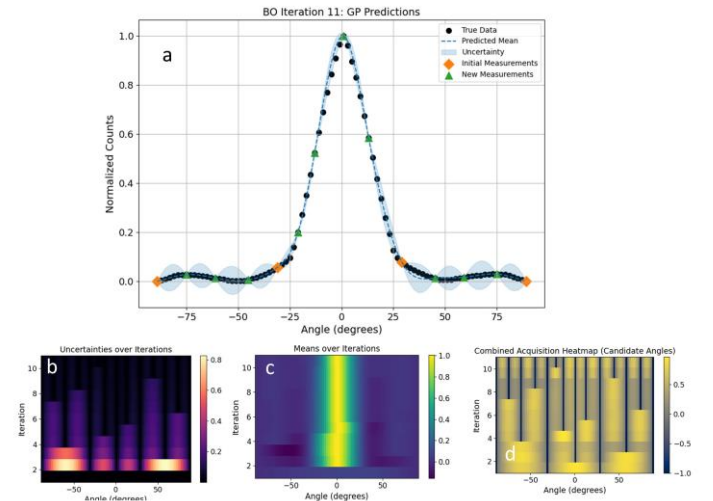


Figure 2: a) Gaussian Process (GP) predicted mean (dashed blue curve) and uncertainty (blue shaded area) after 11 BO iterations, compared against the true dataset (black dots). The orange diamonds represent initial measurements used for model initialization, while green triangles denote measurements acquired based on BO recommendations. The radiation source is centered at  $0^\circ$ , with the scan covering  $180^\circ$  horizontally at  $2^\circ$  intervals; b) Heatmap of GP predictive uncertainty evolution over successive iterations; c) Heatmap of the GP predicted mean radiation distribution over iterations, with the topmost row representing the final estimated distribution shown in (a); d) Combined acquisition function heatmap illustrating candidate angle selection at each iteration. Vertical blue lines indicate the angles selected by the algorithm, starting from the bottom.

Figure 3 presents the results of applying the BO algorithm using data that were obtained at the thermal irradiation port of the MARK II TRIGA research reactor (Jožef Stefan Institute, Slovenia). The experimental setup involved positioning the gimbal at a distance of 3.4 meters from the centre of the beam, with the scan covering a  $180^\circ$  angular range with a  $6^\circ$  angular resolution. Spectral measurements were acquired over a 60-

second integration time producing spectra across 32 angular positions. In Figure 3a, the algorithm is applied to the gross spectral counts, while in Figure 3b, it is applied on the angular response of the 511 keV  $\gamma$ -ray peak associated with electron-positron annihilation. In both cases, the adaptive sampling method demonstrates its capability to recover the true underlying radiation distribution while requiring only a 45 to 50% of the raster scan measurements.

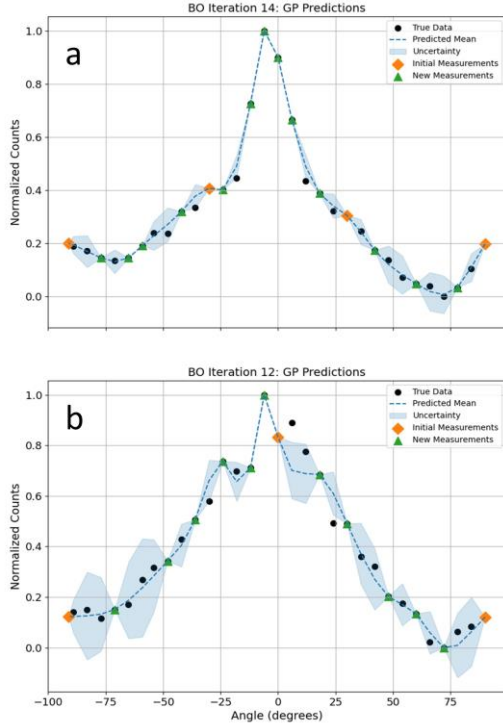


Figure 3: (a–b): In-situ tests conducted at the TRIGA research reactor’s thermal irradiation port, and estimated radiation distributions produced by the BO algorithm on the gross spectra (a) and 511 keV peak (b). In each subplot, the black dots represent the true dataset obtained from full raster scans, the orange diamonds indicate initial measurements, and the green triangles mark the locations selected by the BO algorithm. The dashed blue curve shows the predicted mean distribution, and the shaded blue region represents the prediction uncertainty.

### B. Hardware implementation

Figure 4 depicts results from the practical implementation of the algorithm using the gimbal-mounted collimated  $\text{CeBr}_3$  detector with one and two  $^{137}\text{Cs}$  sources. Each plot presents the BO-driven estimated distribution together with the sampled locations, and associated uncertainty. The distribution obtained from raster scans carried out independently is also depicted. The adaptive approach closely matches the detailed angular radiation distributions obtained separately via the exhaustive raster scans (carried out in  $2^\circ$  intervals) with only a quarter of measurements needed; thus, validating the results obtained in the previous subsection.




Figure 4: Experimental validation of Bayesian Optimization scans using the gimbal-mounted collimated  $\text{CeBr}_3$  detector with (a) a single  $\text{Cs-137}$  radiation source and (b) two  $\text{Cs-137}$  radiation sources. The BO-sampled measurements (solid black circles) and initial sampling points (open black circles) are compared with reference raster scans (red crosses). The Gaussian Process Regression (GPR) predicted mean angular response (solid black line) and its uncertainty bounds (dashed black lines) suggest a near perfect match with the raster scan results.

## IV. DISCUSSION

The results demonstrate that substantial efficiency gains can be achieved using this adaptive sampling approach compared to sequential scanning. The algorithm consistently converges with significantly fewer measurements while yielding accurate angular radiation distributions. This reduction is critical in robotic radiation localization applications, where rapid and accurate assessments are of essence.

On the other hand, there are certain aspects to further address and refine in order to create a more robust implementation of the algorithm. Firstly, the choice of the kernel for the Gaussian Process, although adequate in the controlled scenarios presented here (point sources and beam port conditions), might be suboptimal or a poorly matched prediction tool for scenarios that involve dispersed contamination. Planned experiments involving such non-point source distributions will explore this hypothesis and the potential of implementing dynamic kernel selection.

Secondly, the current formulation of the acquisition function presents potential drawbacks: specifically, when the  $\alpha$  parameter in equation 5 favours the influence of the Upper Confidence Bound (UCB) component, and as it becomes more dominant, the algorithm tends to over-exploit regions of high radiation intensity, resulting in more measurements than necessary. Therefore, a means of monitoring and dampening such behavior in real-time is essential, without hindering the

exploration-exploitation balance. For instance, the geometric characteristics of the collimator, such as its field of view, can be utilized to impose larger angular intervals than  $1^\circ$  if the algorithm gets trapped in overexploitation.

Another aspect that needs addressing is that the aforementioned convergence criteria are satisfied using manually set thresholds. These are inherently heuristic and might not apply in all scenarios, such as the extreme contrast of low and high background between a laboratory equipped with point sources and an operating reactor. Implementing a real-time, automated setting of the thresholds based on observed measurement patterns could further optimize the sampling efficiency.

Finally, further research will focus on testing this algorithm with dispersed sources and subsequently testing its performance when integrated as the main radiation sensor for on-board high level decision-making algorithms, such as radiation-seeking or radiation-avoidance navigation

## V. CONCLUSIONS

This paper proposed a Bayesian Optimization approach that offers a practical and efficient alternative to conventional scanning methods for  $\gamma$ -localization. The results suggest that it can maximize input space coverage without compromising accuracy, and hence it can be a useful instrumentation method for *in-situ* robotic characterization tasks. Future work will focus on enhancing its performance, such that it can be applied in broader and more complex sets of conditions.

## ACKNOWLEDGMENTS

We acknowledge the financial support of the Engineering and Physical Sciences Research Council (EPSRC) via grants no. EP/V026941/1 and EP/X038327/1. Malcolm Joyce acknowledges the receipt of a Wolfson Research Merit Award from the Royal Society.

## REFERENCES

- [1] I. Tsitsimpelis, C. J. Taylor, B. Lennox, and M. J. Joyce, ‘A review of ground-based robotic systems for the characterization of nuclear environments’, *Progress in Nuclear Energy*, vol. 111, 2019, doi: 10.1016/j.pnucene.2018.10.023.
- [2] A. West *et al.*, ‘Use of Gaussian process regression for radiation mapping of a nuclear reactor with a mobile robot’, *Sci Rep*, pp. 1–11, 2021, doi: 10.1038/s41598-021-93474-4.
- [3] B. Nouri Rahmat Abadi *et al.*, ‘CARMA II: A ground vehicle for autonomous surveying of alpha, beta and gamma radiation’, *Front Robot AI*, vol. 10, Mar. 2023, doi: 10.3389/frobt.2023.1137750.
- [4] H. Chapman *et al.*, ‘Robots in nuclear decommissioning Safety Case Considerations for the Use of Robots in Nuclear Decommissioning’, *VGB PowerTech*, vol. 5, no. 2020, 2020.
- [5] H. Chapman, J.-P. Richardson, C. Fairbairn, A. Di Buono, and A. Gale, ‘Best Practice Guidelines and Lessons Learned from Robotic System Deployment in Nuclear Decommissioning’, *ATW International Journal for Nuclear Power*, vol. 68, pp. 53–59, 2023, [Online]. Available: <https://orcid.org/0000-0001-7844-8890>
- [6] A. Chierici, R. Ciolini, A. Malizia, and F. d’Errico, ‘Resource Constrained Electronics and Signal Processing for UAV Radiation Sensors’, in *EPJ Web of Conferences*, A. Lyoussi, F. D’Errico, M. Carette, M. Joyce, R. Hodák, I. Jenčič, P. Le Dù, M. Morichi, S. Pospíšil, C. Reynard-Carette, L. Snoj, and L. Vermeeren, Eds., Nov. 2023, p. 10019. doi: 10.1051/epjconf/202328810019.
- [7] I. Tsitsimpelis *et al.*, ‘Low-density collimators for remote radiation characterisation tasks’, *Nucl Instrum Methods Phys Res A*, vol. 1079, p. 170701, Oct. 2025, doi: 10.1016/j.nima.2025.170701.