# Robotic Neutron Detection for Bayesian Source Localization, Absence Confirmation, and Template Matching

Eric Lepowsky,* Alexander Glaser,* Robert J. Goldston,*†

*Program on Science and Global Security, Princeton University
†Princeton Plasma Physics Laboratory

**Abstract.** In scenarios where no neutron sources or significant changes are declared, the ability to characterize a neutron field and identify possible anomalies could be crucial for verifying compliance to safeguards or arms control agreements. This work demonstrates a unified approach using the N-SpecDir Bot, a Neutron-detecting Spectrally and Directionally sensitive robot, for localization of an anomalous source, absence confirmation, and neutron field template matching. Starting from a simplified particle filter, we introduce several improvements which are applicable to potential applications. We generalize the particle filter to account for unknown background radiation and the possibility of zero sources so that it can simultaneously verify the absence of sources or localize an anomalous source while incorporating information from successive measurements as the N-SpecDir Bot explores its search environment. We also propose a method for adapting the particle filter to template matching with prior measurements by redefining the particles to represent the location and magnitude of deviation from the template. Template acquisition, interpolation, and matching is demonstrated with experiments conducted at Princeton Plasma Physics Laboratory.

## Background

Verification of nuclear safeguards and arms control agreements may require labor-intensive and intrusive onsite inspections. The introduction of robotic inspection capabilities has the potential to bring about an increase in efficiency, effectiveness, and safety.[1] A "robotic inspector" enables remote inspections, wherein a human need not be present at the inspected site, thereby reducing safety and security concerns. Inspection tasks may include a variety of typically human-based measures, including verification of labels and seals, counting of objects, and, of principle interest here, radiation measurements. Leveraging the mobile, and potentially autonomous, nature of a robotic radiation detector, this work aims to provide the ability to characterize a neutron field and identify possible anomalies in scenarios when no significant neutron sources or changes are expected. Herein, we focus on our efforts to demonstrate a unified approach using the N-SpecDir Bot, a Neutron-detecting Spectrally and Directionally sensitive robot, for localization of an anomalous source, confirming the absence of sources, and template matching of neutron fields.

The N-SpecDir Bot (Figure 1) is comprised of three or six boron-coated straw (BCS) detectors azimuthally-distributed within a cylinder of high-density polyethylene (HDPE), which is mounted on an omni-directional robotic platform.[2] The HDPE attenuates neutrons as they traverse through the detector system, generating a neutron flux gradient between the detectors. The HDPE is surrounded by an aluminum and cadmium casing, the latter of which

shields the detectors from an influx of thermalized neutrons. While many reported methods for source localization use only total detected counts, our N-SpecDir Bot is specifically designed to provide directional sensitivity,[3] in addition to gross counts, by utilizing the signals from the independent detectors.[4]

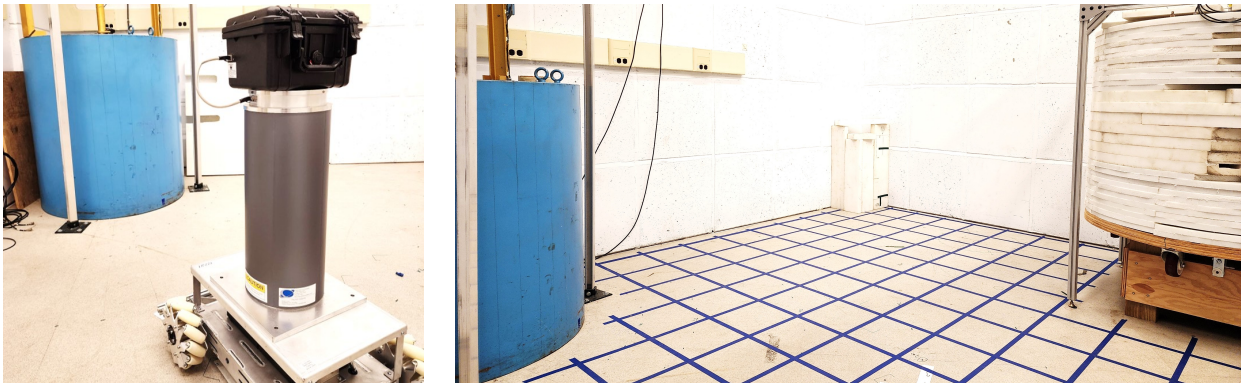

**Figure 1: N-SpecDir Bot in the Calibration Service Laboratory (CASL) at Princeton Plasma Physics Laboratory (PPPL).** Left: The N-SpecDir Bot stands in front of a containerized PuBe source (blue cask) in CASL. Right: The full source range with a 30-cm grid used for positioning during neutron field measurements.

Starting from a simplified particle filter, we introduce several improvements which are applicable to potential applications of interest, from monitoring of the uranium enrichment process[5] to verifying the absence of undeclared warheads. In the following sections, the particle filtering framework is first introduced.[6] We then generalize the particle filter so that it can simultaneously verify the absence of sources or localize an anomalous source. Subsequently, we propose a method for adapting the particle filter to neutron field template matching. An experimental demonstration of template acquisition, interpolation, and matching is then presented before discussing directions for ongoing and future work.

## Source Localization

***Basics of Particle Filtering.*** The particle filtering framework is selected for its ability to incorporate information from successive measurements as the N-SpecDir Bot explores its search environment. Each "particle" represents a hypothesized source and is assigned a weight according to a cumulative likelihood function. The goal is to determine the probability of source hypotheses given the acquired measurements; however, such a probability is difficult to infer directly. Conversely, the probability of a measurement given a source is tractable, and Bayes Rule allows us to use this information to yield the desired result. In the case of particle filtering without re-sampling, at time step $t$, the posterior belief $bel(x_t) = P(x_t|z_{1:t})$ is estimated by updating the particle weights according to Eq. 1, for source hypothesis $x$ and measurement $z$.[7] Without re-sampling, the particles do not evolve in time, so $x_t$ can simply be denoted by $x$, and by constraining the particles to a consistently spaced grid,[8] the particle filter is essentially reduced to grid-based recursive Bayesian filtering.

$$w_t = P(z_t|x_t) \times w_{t-1} \tag{1}$$

A sensor model, $P(z_t|x)$, must be defined to assign weights to the particles, where the weight represents the probability of the corresponding particle yielding the acquired measurements. Assuming independence of measurements, the joint probability of acquiring multiple measurements is given as $\prod_t P(z_t|x)$. However, as the sensor model uncertainty decreases or the number of measurements increases, the joint probability may tend toward zero, so it is advantageous to convert to logarithmic likelihoods. For the probability mass function (PMF, discrete) or density function (PDF, continuous) given by $P(z|x)$, let $f(z|x)$ be the corresponding logarithmic likelihood function (log-likelihood, for brevity); throughout this work, the iteration index, $t$, is neglected for notational simplicity. Since the logarithm of a product is the sum of logarithms, the log-likelihood of measurements taken at different times or positions can simply be summed, as well as measurements of different forms. With the N-SpecDir Bot, we have information on total counts and direction, so $f(z|x) = f(z_c|x_c) + f(z_\theta|x_\theta)$, where $z_c$ and $z_\theta$ are the measured counts and direction, and $x_c$ and $x_\theta$ are the predicted counts and direction, corresponding to the source hypothesis $x$.

***Likelihood Functions for Source Localization.*** The PMF for detecting the total measured counts summed between all detectors, $z_c$, given the hypothesized source counts, $x_c$, is defined in Eq. 2 as the Poisson distribution. The corresponding logarithmic likelihood function is given by Eq. 3, which is further simplified, denoted by the primed function, to neglect terms which only depend on the measurement and are constant for all particles.

$$P(z_c|x_c) = e^{-x_c}\frac{{x_c}^{z_c}}{z_c!} \tag{2}$$

$$f(z_c|x_c) = -\ln(z_c!) + z_c\ln(x_c) - x_c \Longrightarrow f'(z_c|x_c) = z_c\ln(x_c) - x_c \tag{3}$$

The PDF for measuring the direction, $z_\theta$, given the hypothesized direction, $x_\theta$, is defined in Eq. 4 as the von Mises distribution. $x_\theta$ is the apparent direction based on the predicted counts in each detector, which for simple cases with a single source, is simply the direction from the current robot position to the hypothesized source location. $I_0$ in Eq. 4 is the zeroth order modified Bessel function of the first kind. The corresponding logarithmic likelihood function is given by Eq. 5. We assume that the concentration may be taken as the inverse of the variance, $\kappa \approx \frac{1}{\sigma_\theta^2}$. The variance, $\sigma_\theta$, in the directional estimate is computed by propagating the Poisson statistics from the predicted counts. The approximation for the concentration is valid for $\kappa \gg 0$ or, analogously, small $\sigma_\theta^2$. Although the von Mises distribution more closely resembles a uniform distribution for small $\kappa$, we assume the inverse relation between concentration and flattened dispersion holds for the sake of estimation.

$$P(z_\theta|x_\theta,\kappa) = \frac{e^{\kappa\cos(z_\theta - x_\theta)}}{2\pi I_0(\kappa)} \Longrightarrow P(z_\theta|x_\theta) \approx \frac{e^{\frac{\cos(z_\theta - x_\theta)}{\sigma_\theta^2}}}{2\pi I_0\left(\frac{1}{\sigma_\theta^2}\right)} \tag{4}$$

$$f(z_\theta|x_\theta) \approx \frac{\cos(z_\theta - x_\theta)}{\sigma_\theta^2} - \ln\left(2\pi I_0\left(\frac{1}{\sigma_\theta^2}\right)\right) \tag{5}$$

## Absence Confirmation

***Predicting an Unknown Background Intensity.*** To perform absence confirmation, we make $z_c$ equal to the total detected counts and define $x_c$ as the predicted counts from the hypothesized source, $S$, plus the counts from the predicted or hypothesized background, $b$. As a result, the predicted mean counts are $x_c = S+b$ with corresponding variance $\sigma_c^2 = S+b$. By computing the log-likelihood of Eq. 3 based on the total counts, as opposed to background-subtracted counts, the background may be varied as an unknown parameter. The predicted source counts, $S$, are a function of three unknown parameters: the $x$-position of the source, the $y$-position of the source, and the intensity of the source. The background, $b$, is then the fourth parameter. So, the distribution of particles forms a four-dimensional space, where each particle is of the form: $[x, y, I, b]$. This implicitly assumes that the background is spatially and temporally uniform. For the case of known background, which may include spatial or temporal variation, $b$ may be set as a known variable, which collapses the distribution of particles down to a three-dimensional space of solely the source parameters.

***Introducing Zero-Intensity Sources.*** Defining $x_c$ as the total predicted counts from both the source and the background also enables the inclusion of "zero-intensity source" particles in the set of possible hypothesized sources.[9] In essence, a "zero-intensity source" particle represents the hypothesis that no source is present at the given location. However, for the sake of applying the same log-likelihood function for all hypotheses, this must be represented by a notional source of zero intensity.[10] By including zero-intensity particles in the distribution of hypotheses, the particle filter may now simultaneously localize a source or confirm the absence of sources; the relative probability of each scenario, and the selection of which task dominates, will inherently be computed by the Bayesian statistics of the particle filter without *a priori* knowledge of the required verification task.

## Template Matching

***Redefining the Particles with a Template.*** While the previous section expanded the particle filter from source localization to absence confirmation, the definition of the particles still implicitly assumes a single source, or no source, to be present. We now wish to consider cases where multiple sources are known to be present.[11] Assume we have a template representation of the neutron field in the search area. If all we were interested in answering is whether the measured counts match the template, we could determine the predicted counts from the template and directly calculate how well the measured counts match. However, if instead we're interested in knowing the amount and location of deviation from the template, the particle filter can provide our answer, and do so without re-acquiring the entire template. This is still a problem of either source localization (finding the largest deviation) or absence confirmation (confirming that the deviation from the template is due to statistical noise).

For each position, we store three parameters in the template, corresponding to the three coefficients in the cosine model (Eq. 6) which characterizes the directionality of the N-SpecDir Bot: $A$, $B$, and $\theta_0$, where $\theta_0$ is the apparent direction and $\theta_i$ is the orientation of the $i^{\text{th}}$ detector. Given this template, we can predict the counts in each of the $i \in [1, ..., N]$

detectors as $T_i$, with the total template counts given by $T = \sum_i^N T_i$. The corresponding variance in the predicted template counts is then $\sigma_{T_i}^2$ and $\sigma_T^2$. Similar notation is used for the predicted source counts, $S = \sum_i^N S_i$.

$$T_i = A + B\cos\left(\theta_0 - \theta_i\right) \tag{6}$$

We continue to use Eq. 3 as the log-likelihood function for counts. The total hypothesized counts $x_c$ include the template, $T$, plus/minus the deviation from the template, $S$, and background, $b$, so $x_c = T + S + b$ with corresponding variance $\sigma_c^2 = T + S + b$.[12] The particles now represent the predicted location and magnitude of deviation from the template. The same "zero-intensity source" particle as before now is representative of a match to the template at the given $(x, y)$ position. Unlike before, the source counts predicted by the particles can be negative in the case of a missing source. Even for a negative deviation, the sum of the template and the deviation will still be non-negative, since the missing source cannot exceed the preexisting one; equivalently, $S$ has a lower bound of $-T$. Intuitively, $S > 0$, represents the case of an additional source being present, superimposed on top of the template, whereas $S < 0$ represents a partially or completely missing source.

***Reconsidering Directionality and Uncertainty.*** We must also consider how the directional estimate and uncertainty propagation are affected by the template-matching scheme. The cosine model should be applied to the *template-subtracted* counts, $\bar{z_c} = z_c - T$, in order to find the apparent direction, $z_\theta$, toward the largest deviation from the template, rather than measuring the direction toward the strongest source in the room. In this manner, the apparent direction is not toward the predominant source, but rather the predominant difference from the template. Analogously, $x_\theta$ will be based on the deviation from the template, $S + b$, rather than on the total counts, $T + S + b$. This can alternatively be thought of as the template-subtracted predicted counts, or $\bar{x_c} = x_c - T$.

Uncertainty must be propagated from the hypothesized counts to the predicted direction. While the variance in the predicted template counts is given by $\sigma_T^2 = T$, the uncertainty in the template itself is given by $\sigma_t^2$, which may not be equal. If the template were acquired with long-duration measurements, $\sigma_t^2 \to 0$, whereas for finite measurement time, $\sigma_t^2 \neq 0$. If the measurement time for acquiring the template is the same as the measurement time used during the subsequent search, then $\sigma_t^2 = \sigma_T^2$. Additionally, if the template is background-subtracted, there would be an additional factor in the uncertainty. To capture all these possibilities, we use $\sigma_t^2$ for generality. The calculation is summarized by Eq. 7, where the variance of the counts in each detector being propagated is given by $T_i + S_i + b_i + \sigma_{t_i}^2$.[13] Note that the original template-free particle filter can be recovered by simply setting $T = \sigma_t^2 = 0$.

$$\sigma_\theta = \sqrt{\sum_{i=1}^{N} \left(\frac{\partial\theta_0}{\partial\bar{x}_{c_i}}\right)^2 \sigma_{\bar{x}_{c_i}}^2} = \sqrt{\sum_{i=1}^{N} \left(\frac{\partial\theta_0}{\partial\left(S_i + b_i\right)}\right)^2 \left(T_i + S_i + b_i + \sigma_{t_i}^2\right)} \tag{7}$$

## Experimental Demonstration

***Neutron Field Measurements and Datasets.*** An experimental campaign was conducted at the Calibration and Service Laboratory (CASL) at Princeton Plasma Physics Laboratory (PPPL) to acquire various neutron field measurements using the N-SpecDir Bot. The CASL provides a 3.6 by 4.2 meter shielded environment (Figure 1) with a fixed plutonium-beryllium (PuBe) source, which was used in its cask to yield count rates ranging from approximately 10 to 120 cps. An additional portable PuBe check source was used to create multi-source neutron fields. Both sources emitted on the order of 1e4–1e5 n/s, through their containers, at the time of measurement. For each neutron field measurement, a 30-cm grid was used to acquire 50–60 five-minute measurements in a checkerboard pattern. The baseline field (Figure 2) consisted of only the containerized PuBe source; in addition to the checkerboard, a set of 22 complementary measurements were acquired for testing and validation.

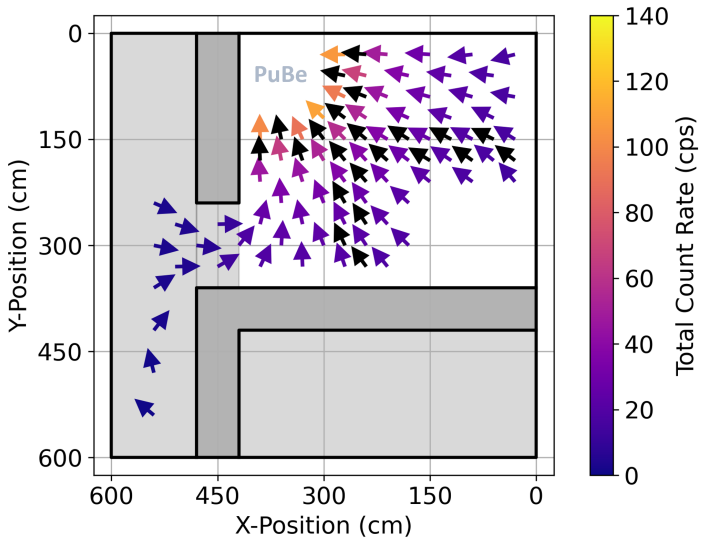


**Figure 2: Baseline neutron field.** The arrows point in the measured direction, $\theta_0$, while the coloring indicates the count rate, $T$, summed between the six detectors of the N-SpecDir Bot. The dark shaded regions indicate shield walls. The measured direction forms a flow field which consistently points into the room and toward the containerized PuBe source located near (370 cm, 50 cm). The complementary data points for validation are shown in black. The non-shaded region is utilized for subsequent neutron field measurements.

Three multi-source neutron field measurements (Figure 3) were acquired using the check source. First, the check source was placed in a corner of the room, away from the containerized PuBe source. Second, the check source in the corner was shielded with 2-in thick polyethylene blocks. Third, the check source was re-positioned in the center of the room, such that measurements could be acquired from all sides of the source.

For the purpose of template matching, we wish to compare any subsequently acquired measurements to any template, which means the measurement points may not be the same. To accommodate this discrepancy, the stored template(s) must be interpolated, such that the expected counts and direction can be referenced at any accessible point in the environment.

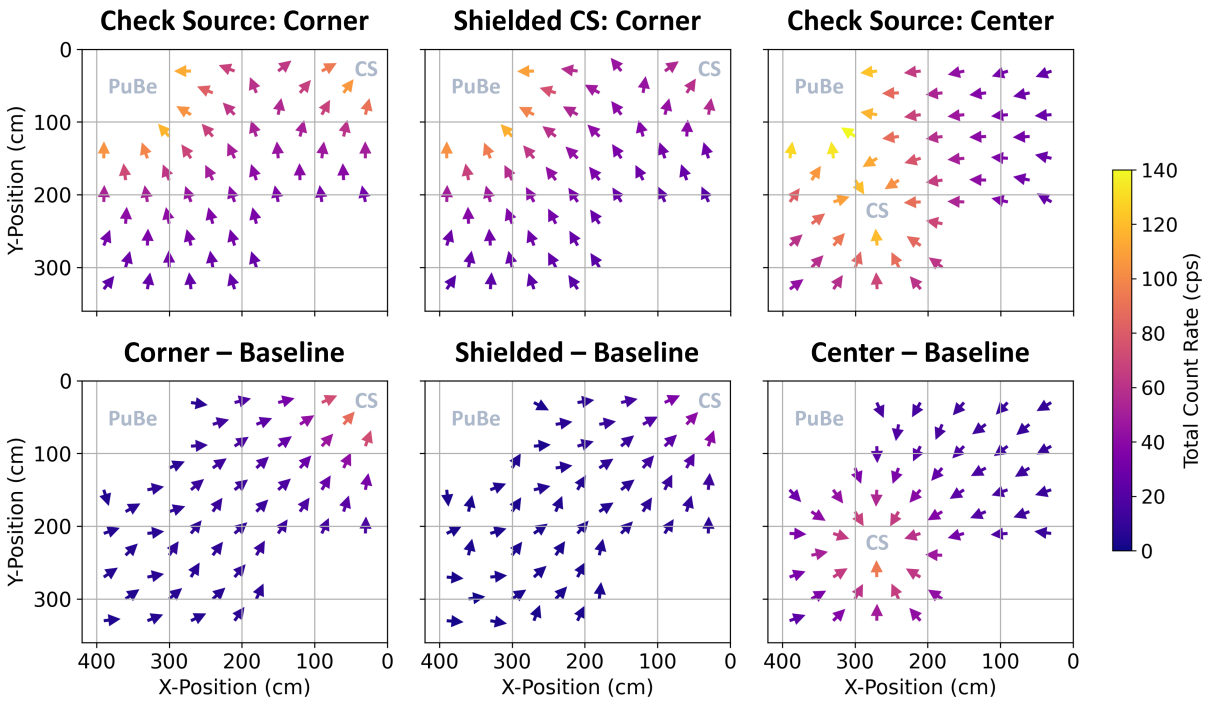


**Figure 3: Multi-source neutron fields and differences.** The top row consists of the neutron field measurements in CASL using the containerized PuBe source and neutron check source (CS). The measured direction is skewed by the check source, causing saddle points to form in the flow field. The bottom row is the result of subtracting the baseline field from the top row. Left: Check source located at (10 cm, 2 cm). Middle: Check source located at (10 cm, 2 cm) behind 2-in thick polyethylene shielding. Right: Check source located at (270 cm, 225 cm).

In order to wrap smoothly around points where the apparent direction rapidly changes, such as around a source or at the saddle point between two sources, $\cos(\theta_0)$ and $\sin(\theta_0)$ are interpolated independently. A variety of Python-based two-dimensional interpolation methods were compared to down-select the most accurate method for this application and dataset. Count rate and direction error were defined as $|T_{true} - T_{interp}|/T_{true}$ and $|\Delta\theta_0|/180$, respectively, where $\Delta\theta_0$ is the minimum angular distance between $\theta_{0,true}$ and $\theta_{0,interp}$ in degrees. The results are summarized in Table 1, for which the baseline field was used for interpolation and compared to the complementary validation test points.[14]

| | Nearest Neighbor | Linear | Cubic | **Bivariate Spline** | RBF Cubic | RBF Quintic |
|---|---|---|---|---|---|---|
| Count Rate Error (%) | 16.01 | 3.85 | 2.81 | **1.91** | 2.79 | 2.23 |
| Direction Error (%) | 3.54 | 1.19 | 1.17 | **1.17** | 1.13 | 1.18 |

**Table 1: Interpolation error for baseline neutron field template.** Interpolation error averaged over 22 validation points for the baseline neutron field template. The two best-performing kernels (of eight possible built-in options) are included for radial basis function (RBF) interpolation. The method with the lowest overall error (bivariate spline) is highlighted in bold and subsequently used throughout this work.

***Template Matching with Experimental Data.*** The particle filter operates in the "difference" between acquired neutron fields, as visualized in Figure 3. Using the baseline neutron field from Figure 2 as the template and any of the multi-source fields as the measurement results in a source localization problem (localizing the addition of the check source). The inverse scenarios also entail source localization, but with negative source intensity (localizing the subtraction or removal of the check source). If the template and measurement fields match, this is the case of absence confirmation (confirming the absence of deviations from the template). Figure 4 demonstrates the application of the modified particle filter, using the experimental data for both the template and measurement fields, to perform template matching; the particle filter successfully accomplishes both source localization (identification of the check source) and absence confirmation (confirming a match to the template), without *a priori* knowledge of the verification task. In this implementation, only the total count information is used; incorporating the directional information is ongoing work.

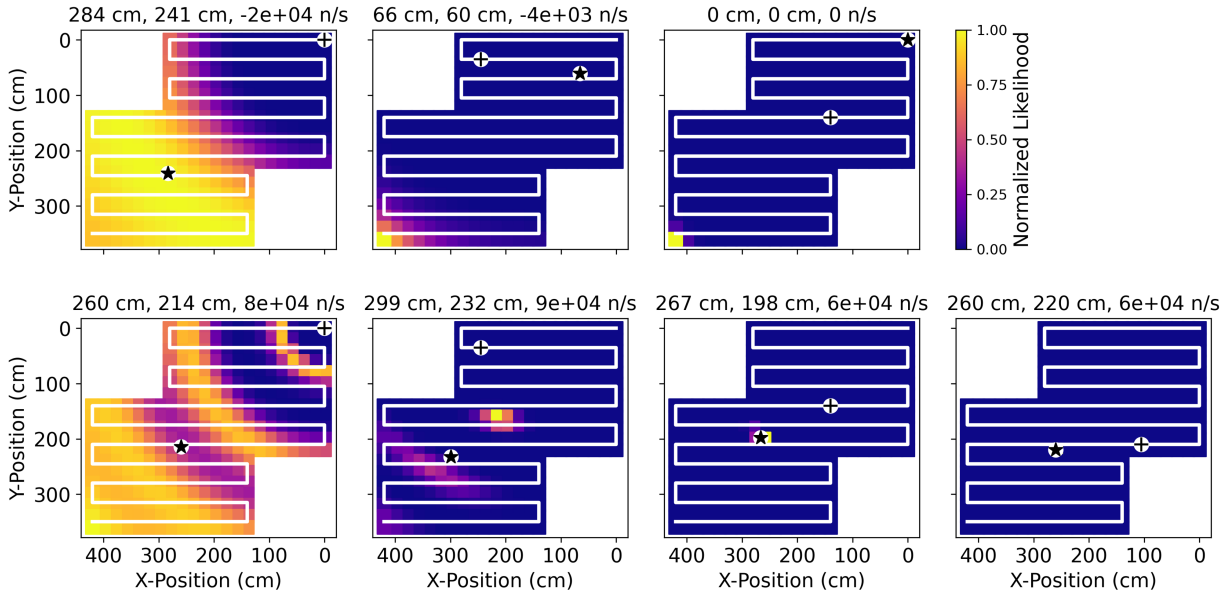


**Figure 4: Particle filter snapshots for source localization and absence confirmation scenarios.** The top row demonstrates absence confirmation, where both the template and the measured neutron field are the "baseline" field. The bottom row is an example of source localization, where the template is the "baseline" field and the measured neutron field is the "center" field (see Figure 3). For both scenarios, a 35-cm resolution raster pattern is followed with 10-s measurements. The panels, from left to right, visualize iterations 0, 10, 40, and 65. Above each panel is the weighted center of mass of particles, which is indicated by the black star. The current robot position is indicated by the black crosshair. The last panel for each row demonstrates a fully converged result.

The measurement time at each dwell point and the step size between dwell points, or equivalently the resolution of the raster pattern, are tunable parameters in the presented implementation. Table 2 summarizes the convergence results for using the particle filter with template matching to achieve absence confirmation, i.e., confirming that the measured neutron field matches the template field. Three raster pattern resolutions were tested for different dwell times; three trials were conducted for each of the four neutron field datasets. The results

demonstrate a general preference for spatial coverage over dwell time, as the total measurement time is lower toward the lower-left corner of the table (greater number of measurements, larger step size, shorter dwell time). Another consideration to be explored in future work is the effect of template sparsity on the convergence speed and accuracy of the particle filter.

| | 1 s | 5 s | 10 s | 30 s | 60 s | 120 s | 300 s |
|---|---|---|---|---|---|---|---|
| 35 cm | 96 (96 s) | 58 (290 s) | 36 (360 s) | 16 (480 s) | 8 (480 s) | 5 (600 s) | 3 (900 s) |
| 50 cm | N/C | 35 (175 s) | 28 (280 s) | 15 (450 s) | 9 (540 s) | 5 (600 s) | 3 (900 s) |
| 80 cm | N/C | 21 (105 s) | 17 (170 s) | 12 (360 s) | 8 (480 s) | 4 (480 s) | 3 (900 s) |

**Table 2: Average number of steps and total measurement time to achieve absence confirmation.** For each measurement time and step size, the number of steps is averaged over three independent trials of each of the four neutron field datasets. For each trial, the first iteration when $\log_{10}(\bar{I}) < -2$ was recorded as reaching convergence, where $\bar{I}$ is the weighted average of the particle source intensities (n/s); the result is marked not-converged (N/C) if absence was not confirmed by the end of the raster pattern. The corresponding total measurement time, neglecting time traversing between dwell points, is reported in parenthesis.

## Discussion & Future Work

This work advances toward a mobile and intelligent platform for performing remote, or potentially autonomous, inspections in support of safeguards and treaty verification with a robotic neutron detector. For scenarios when no significant neutron sources or changes are expected, we propose a unified approach for localizing an anomalous source, confirming the absence of sources, and performing neutron field template matching. We demonstrate the adapted particle filtering framework on experimental measurements of various neutron fields.

For all variants of the particle filtering approach, we can simplify the distribution of particles by incorporating physical limitations on the placement of potential sources, since there cannot be a source co-located with the robot in open space. If a map of the environment is known, the particles can simply be initialized only where there are substantial objects capable of containing a neutron source. If the environment is unknown, alternative methods for exploration can be implemented for eliminating infeasible particles, such as including omni-directional distance measurements using LiDAR (Light Detection and Ranging). The motion planning scheme can also be optimized for efficiency. In the presented work, a raster pattern is constructed based on the predefined accessible space to guarantee coverage of the search environment. For a more adaptive, and likely more efficient approach, active sensing should be employed such that the current, and previous, measurements inform the next best measurement location. There are several reported methods for performing active sensing; however, many source localization algorithms are inherently greedy and would not be well-suited for absence confirmation which requires more exploration. The presented improvements and proposed extensions help to mature particle filtering with the N-SpecDir Bot from a canonical source localization approach to an application-relevant tool.

## Acknowledgements

The experimental demonstration presented in this work would not have been possible without the support of numerous researchers, staff, and the Health Physics team from Princeton Plasma Physics Laboratory. Eric Lepowsky's contributions to this project have been supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2039656. This work was partly supported by the Consortium for Monitoring, Technology, and Verification under Department of Energy National Nuclear Security Administration award number DE-NA0003920.

## Endnotes

[1]Similar motivation was provided by the International Atomic Energy Agency Robotics Challenge: K. Robertson, R. Stohr, A. Elfes, P. Flick, A. Sokolov, D. Finker, and C. Everton, "The IAEA Robotics Challenge – Demonstrating Robots for Safeguards Inspections," IAEA Symposium on International Safeguards: Building Future Safeguards Capabilities, IAEA-CN-267/215, 2018.

[2]The detection system (BCS detectors, HDPE moderator, casing) was provided by Proportional Technologies Inc. and the robotic platform was from SuperDroid Robots, Inc.

[3]In addition to directional sensitivity, the configuration of the detection system also provides spectral sensitivity. A softer neutron source will result in a greater front-to-back ratio, i.e., the ratio of the counts detected in the front detector divided by the counts detected in the back detector(s). In principle, if the position of a source is well-characterized, the front-to-back ratio may be indicative of the energy of the source.

[4]A full description and analysis of the directionality is reported in E. Lepowsky, M. Kütt, S. Aslam, H. Fetsch, S. Snell, A. Glaser, R. J. Goldston, "Experimental Demonstration and Modeling of a Robotic Neutron Detector with Spectral and Directional Sensitivity for Treaty Verification," Nuclear Instruments and Methods in Physics Research A, 1041, 2022.

[5]More in depth discussion of potential safeguards applications is provided in R. J. Goldston, A. Glaser, M. Kütt, P. Landgren, N. E. Leonard, "Autonomous Mobile Directionally and Spectrally Sensitive Neutron Detectors," IAEA Symposium on International Safeguards, 2018.

[6]Preliminary concepts were previously presented at the 2021 INMM & ESARDA Joint Virtual Annual Meeting: E. Lepowsky, A. Glaser, R. J. Goldston, M. Kütt, "Toward Autonomous Robotic Inspections of Nuclear Facilities Using Directionally-Sensitive Neutron Detectors."

[7]Common practice is to dynamically re-sample the particles according to the posterior distribution to improve the calculation efficiency. However, for the scenario presented here, we omit re-sampling for simplicity and transparency.

[8]While the grid for source position follows linear spacing, linear or logarithmic spacing may be used in other particle dimensions depending on the range of possible source intensities.

[9]The simplified Poisson log-likelihood yields an infinite solution if $x_c$ is equal to zero; this is only possible if there is no source and no background, which is a highly unlikely and unrealistic scenario. If $x_c$ was, instead, defined as solely the predicted source counts and $z_c$ was defined as the measured background-subtracted counts, then it would be conceivable to have no source present, thereby yielding an infinite log-likelihood.

[10]This is implemented by either concatenating a set of zero-intensity particles (one particle for each background intensity) or by placing zero-intensity particles at every possible $(x, y)$ position and for every possible background intensity. The implications of the chosen method are handled in the post-processing of the particle weights.

[11]There are reported methods for adapting the particle filter to perform localization of multiple sources. Importantly, the fundamental approach is unchanged, and only the distribution of particles is altered. In this work, we are still interested in the localization of a single source *among* other sources present.

[12]In expressing the predicted counts as $T + S + b$, we assume that the template itself contains background-subtracted values, such that the total counts do not include the factor of background twice.

[13]The expressed form for the template-subtracted variance is made clear when we consider that we are propagating the variance of the template-subtracted predicted counts, $\bar{x_c} = x_c - T$, which would be $\sigma_c^2 + \sigma_t^2 = T + S + b + \sigma_t^2$.

[14]Further interpolation testing may include parameter tuning, considerations for local versus global interpolation, and assessing the effect of template sparsity on interpolation accuracy. Default parameters and all data points (i.e., global interpolation) were used for this preliminary investigation.