Proceedings of the INMM & ESARDA Joint Virtual Annual Meeting
August 23–26 & August 30–September 1, 2021

Toward Autonomous Robotic Inspections of Nuclear Facilities Using Directionally-Sensitive Neutron Detectors

Eric Lepowsky[1], Alexander Glaser[1], Robert J. Goldston[1,2], Moritz Kütt[3]

[1]*Program on Science and Global Security, Princeton University, Princeton, NJ, USA*
[2]*Princeton Plasma Physics Laboratory, Princeton, NJ, USA*
[3]*Institute for Peace Research and Security Policy at the University of Hamburg, Germany*

Abstract. Whether safeguarding Iranian uranium enrichment facilities, denuclearizing North Korea, or verifying limits on U.S. and Russian arsenals, nuclear safeguards and arms-control traditionally require intrusive on-site inspections to perform verification tasks. In such applications, the ability to localize a radioactive source is imperative for identifying anomalies when no significant neutron emitters are expected or declared to be present, such as for confirming the absence of clandestine withdrawal stations in the centrifuge hall of a gas-centrifuge enrichment plant or undeclared warheads in a storage facility. We are interested in the role of autonomous mobile robots, which, if designed properly, may be more effective and efficient and less intrusive than their human counterparts. Toward developing such a capability, we have constructed an "Inspector Bot," comprised of three boron-coated straw detectors azimuthally distributed within a cylinder of high-density polyethylene, which is mounted on an omni-directional robotic platform. While many reported methods for source localization use only total detected counts, our Inspector Bot is specifically designed to provide directional and spectral sensitivity, in addition to gross counts, by utilizing the signals from the three detectors. The detection system has been extensively characterized by MCNP modeling, which has been benchmarked to experiments conducted at the Princeton Plasma Physics Laboratory. For source localization using our Inspector Bot, we utilize a simple system of equations which, with the three detectors, is solved to estimate the direction to the source. We finally apply the result of the directional model in the framework of a particle filter.

Introduction

Nuclear safeguards and arms-control traditionally require intrusive on-site inspections to perform verification tasks. While hand-held radiation measurements are a useful tool, they can be costly or undesirable for routine inspections. If designed properly, autonomous robotic inspectors may be able to perform radiation measurements less intrusively and more effectively, efficiently, and safely than current human-based methods.[1] By replacing or supplementing human inspectors with autonomous mobile robots, sensitive information – such as facility design, warhead design, or radiation measurements – need not cross facility walls. In applications where no significant neutron emitters are expected or declared to be present, such as confirming the absence of clandestine withdrawal stations in the centrifuge hall of a gas-centrifuge enrichment plant or undeclared warheads in a storage facility, the ability to localize a radioactive source is imperative for identifying anomalies. In the former example, elevated neutron count rates in cascade halls, driven by (alpha, n) reactions on fluorine, could reveal the presence of hidden uranium hexafluoride containers.[2] In the latter case, we could consider scenarios pertinent to potential future arms-control treaties which may require verifying the absence of treaty-accountable objects.

To examine the potential of autonomous, mobile, directionally-sensitive neutron detectors to perform such verification tasks, we have constructed an "Inspector Bot" comprised of three (or more [3]) boron-coated straw (BCS) detectors azimuthally-distributed within a cylinder of high-density polyethylene (HDPE), which is mounted on an omni-directional robotic platform.[4] The HDPE attenuates neutrons as they traverse through the detector system, generating a neutron flux gradient between the detectors. The detectors and HDPE are also surrounded by an aluminum and cadmium casing, the latter of which attenuate neutrons below 0.5 eV, shielding the detectors from an influx of thermal neutrons which would otherwise increase the total detected counts and effectively diminish the relative difference in the neutrons captured within each detector. While many reported methods for source localization use only total detected counts, our Inspector Bot is specifically designed to provide directional sensitivity,[5] in addition to gross counts, by utilizing the signals from the three detectors; this functionality relies on the difference between the signals.

In the following section, we present a detailed, experimentally-calibrated MCNP model of the Inspector Bot. Using both experimental and simulated data, we then examine the directional sensitivity of the detection system. We conclude by demonstrating how the directional sensitivity may be leveraged for localizing an unknown neutron source.

Model Calibration

The detection system has been extensively characterized by MCNP modeling,[6] which has been benchmarked to experiments conducted at the Princeton Plasma Physics Laboratory (PPPL). A Cf-252 source (estimated 1.8×10^7 n/s at time of measurement) was used to calibrate and test the detection system. Figure 1 depicts the Inspector Bot on-site at PPPL,

along with the MCNP models of the Bot and the Cf-252 source and storage cask. A series of measurements were performed with various configurations within the TFTR Test Cell at PPPL. The Test Cell is a large (34 m by 45 m floor, 16 m ceiling) concrete room. For calibrating the MCNP model, we used a set of measurements acquired for five-minutes with the Bot facing directly toward the Cf-252 source at three increasing standoff distances. For each measurement, the counts in each detector were simultaneously recorded.

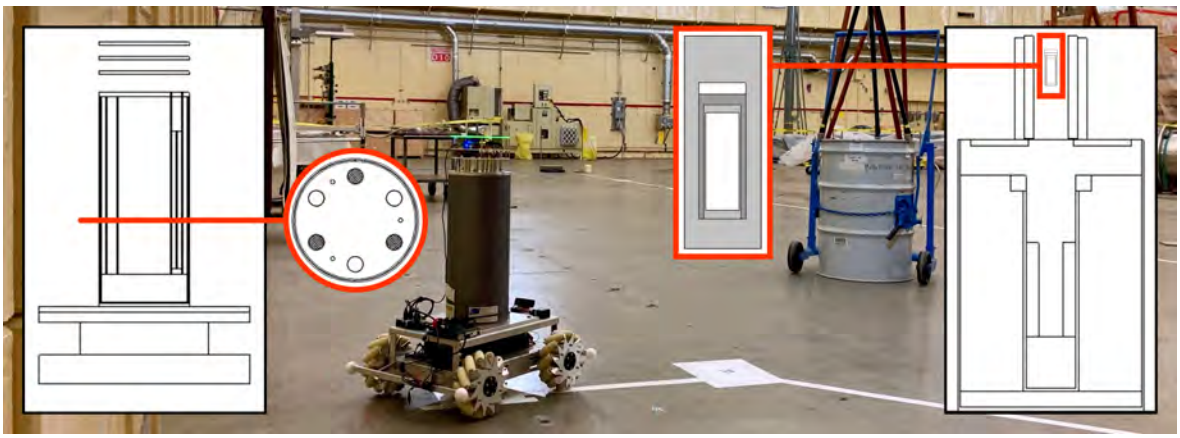

Figure 1: Inspector Bot and Cf-252 source cask located in the TFTR Test Cell at PPPL. To the left of the robot is a slice of the MCNP model (omni-directional wheels are included but out of the image plane); inset is a cross-sectional view of the BCS detectors (1" diameter) and HDPE moderator (8" diameter). The current Inspector Bot has three BCS detectors with room to accommodate six. To the right of the source cask is a cross-sectional view of the MCNP model; inset is a zoomed-in view of the source capsule model, including the stainless steel inner and outer capsules and polymer filling.

A comparison of the MCNP model and the experimental results is visualized in Figure 2, where both the detected count rate and the front-to-back ratio, defined as the ratio of the counts detected in the front detector divided by the average of the two back detectors, are plotted as a function of the standoff distance between the central axes of the Cf-252 and the Inspector Bot. The process of fine-tuning the model to minimize the discrepancy with the experiment revealed a few important details regarding the sensitivity of the detection system.

It is evident in the experimental data that the detected counts do not follow the traditional inverse-square law. Rather, the counts fall off following a shallower slope with respect to distance. We attribute this trend to the effect of the surrounding environment. While the concrete floor was the predominant factor, MCNP modeling revealed that the walls and ceiling also contributed to achieving a model which accurately reproduced the experimental results, even though the nearest wall was 12 m from the Cf-252. Another very important detail was the source capsule. As opposed to using a simple point source in the MCNP model, we include a detailed model of the source capsule, including the stainless steel outer capsule, stainless steel inner capsule, and polymer filling which holds the inner capsule in place.[7] Although the presence of the capsule had a small effect on the absolute counts, the capsule was a key factor in achieving good agreement with the experimental front-to-back ratio. The source spectrum

affects the front-to-back ratio, as well. The Cf-252 spectrum is commonly represented by the Watt spectrum with probability density function $p(E) = C\exp(-E/a)\sqrt{\sinh(bE)}$. While in this work we use the Cf-252 Watt spectrum which corresponds to that found in the ENDF/B-VIII.0 database (a = 1.1800, b = 1.03419), we also note that a harder alternative spectrum (a = 1.025, b = 2.926) results in a lower front-to-back ratio. This observation is indicative of the spectral sensitivity of the detection system. For instance, when a 6 cm thick polyethylene cylinder is positioned around the source capsule, the front-to-back ratio for the 5 m standoff distance increases from 2.48 for the bare source to 3.43 for the moderated source.[8]

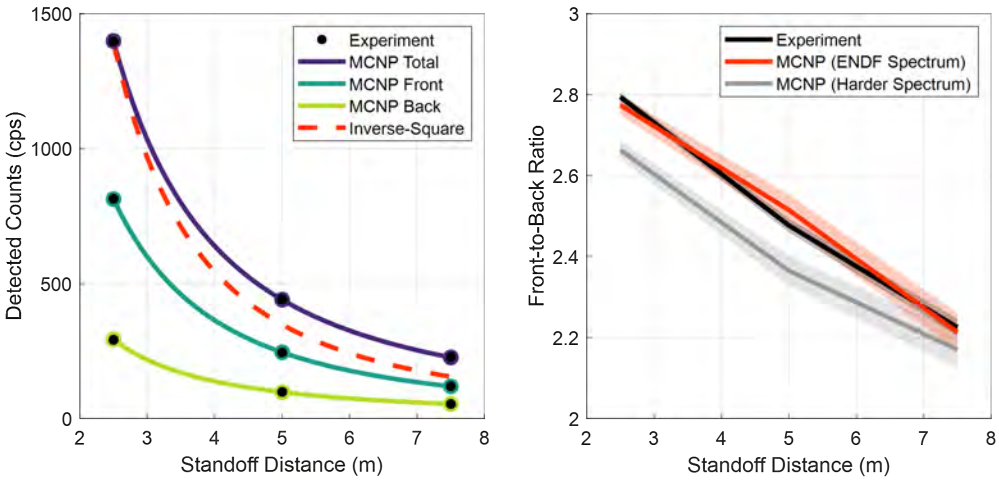

Figure 2: Calibration of MCNP model to experimental measurements. Left: Detected counts in the front detector, average counts of the back two detectors, and total counts of all three detectors as a function of standoff distance for the robot pointed directly toward the source. There is excellent agreement between experimental and simulated data points. Also evident is that the counts do not follow a traditional inverse-square law, attributed to the effect of the surrounding environment (concrete floor, ceiling, and walls). **Right:** Front-to-back ratio (calculated as the front divided by the average of the back) as a function of standoff distance. One-sigma statistical error bars are included for the experiment, as well as the one-sigma interval from MCNP calculations. The Cf-252 Watt spectrum corresponding to the ENDF/B-VIII.0 database (a = 1.1800, b = 1.03419) yields the best fit to the experimental data and is used throughout this work. The effect of a harder Watt spectrum (a = 1.025, b = 2.926) is also shown.

Directional Sensitivity

The detection system of the Inspector Bot was designed to provide directional and spectral sensitivity, the former of which is of principle interest here. In Figure 3, we compare experimental and simulated data as the Inspector Bot is rotated in place. In open space, the shape of the counts-versus-angle curve produced as the detectors sweep through a rotation closely resembles that of a cosine wave; the detected counts are highest for the detector directly facing the source, lowest for a detector located 180° from the forward position, and symmetric about the forward direction. Due to the cosine-like shape of the counts-versus-angle curve,

we represent the angular dependence of the detected counts by Eq. 1, which forms a system of three equations with three unknowns: A, B, and the direction toward the source θ_0, where S_i is the detected counts and θ_i is the relative rotation of the i^{th} detector for $i \in [1, 2, 3]$.

$$S_i = A + B\cos(\theta_i - \theta_0) \tag{1}$$

A single set of measurements consists of the counts detected in each of the three detectors, so the system of equations can be solved by simple least-squares fitting. We note that A is approximately equal to the average detected counts per detector and is weakly dependent on orientation. Furthermore, we choose the bound $\theta_0 \in [-60, 60]$ relative to the detector with the highest counts, denoted here in degrees; in this manner, this simple model enables the Inspector Bot to differentiate between sources positioned at any arbitrary angle.

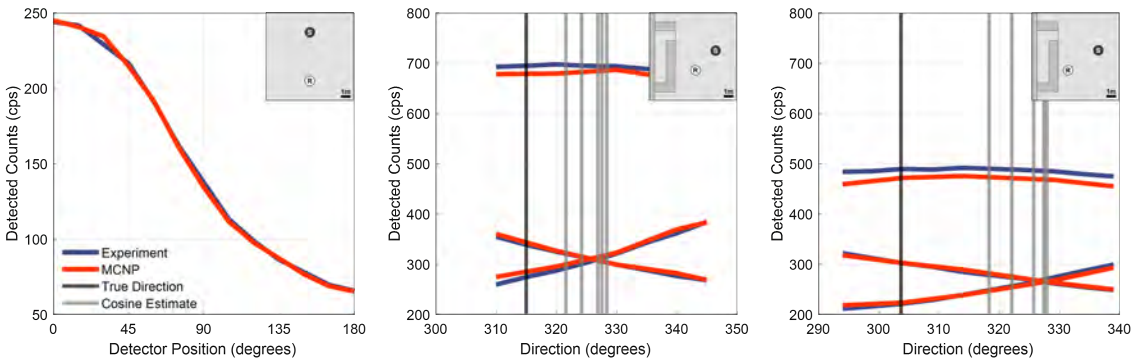

Figure 3: Comparison of experimental and simulated measurements as a function of detector rotation. For each of the cases shown, the inset schematic in the upper-right shows the configuration of the robot (R, white) and source (S, black). **Left:** Detected counts in a single detector as the robot is rotated in open space. **Center:** Detected counts in each of the three detectors. The robot is 2 m away from a concrete wall with the source located at 315° (indicated by the black vertical line). Each curve consists of five measurements; for each measurement, the estimated direction toward the source is indicated by a gray vertical line. **Right:** An additional near-wall case where the robot is moved 1 m closer to the wall with the source stationary, now located at 303.7°. For both near-wall cases, neutrons reflecting off the wall shift the detected direction.

In addition to the detected counts in open space, Figure 3 also shows two "near-wall" cases in which the Cf-252 source and the Inspector Bot are located in close proximity to a concrete wall. The neutrons reflected by the wall effectively shift the apparent direction to the source. For each of these cases, the curve consists of five discrete measurements, where each measurement is itself the detected counts in each of the three detectors. The result of the cosine model applied to each measurement is shown. This demonstrates that with only three data points, the cosine model is able to predict the angle which maximizes the counts in the forward-most detector.[9] We also see the effect of the neutron-reflecting concrete wall, in that the detected direction is approximately 10° to 25° away from the true direction. To correct for this limitation of the cosine model, we are working toward the next iteration of the Inspector Bot, which will contain six neutron detectors instead of three. With the added detectors, we

increase the curve fitting capability to six parameters, and we can unlock higher order terms in the model. From preliminary investigation, we have evidence which may demonstrate the ability to compensate for the effect of large neutron reflectors. In addition to potentially improving the directional sensitivity of the Inspector Bot, the six detectors also provide nearly twice the total count rate, thereby lowering the uncertainty in the total detected counts.

With the current three-detector configuration of the Inspector Bot, we also examine the uncertainty in the directional estimate. The statistical uncertainty in the detected counts is propagated to the direction estimate by $\sigma_{\theta_0} = \sqrt{\sum_{i=1}^{3} \left(\frac{\partial \theta_0}{\partial S_i}\right)^2 \sigma_{S_i}^2}$. For Poisson distributed counts, the variance is equal to the counts, $\sigma_{S_i}^2 = S_i$, and we numerically approximate the partial derivative. Uncertainty propagation was evaluated for varying signal-to-background ratio and increasing standoff distance and is visualized in Figure 4, where we have conservatively fit a linear relationship. The result is remarkable: for standoff distances up to 10 m, source strengths of 10^6 to 10^9 n/s, and signal-to-background ratios of 10^{-4} to 10^4, the directional uncertainty in degrees is never more than twice the percent uncertainty in detected counts.

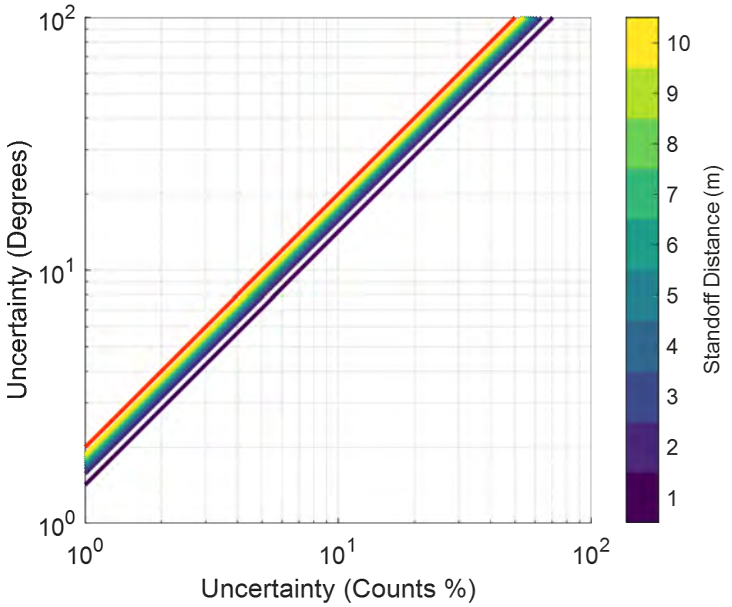

Figure 4: Propagated uncertainty in the Inspector Bot's directional estimate. A conservatively drawn linear relationship between the directional uncertainty in degrees (y-axis) and the percent uncertainty in the detected counts relative to background-subtracted counts for varied signal-to-background (x-axis). As a reference, the orange line is $y = 2x$, which demonstrates that the directional estimate is pessimistically twice as uncertain as the detected counts for the range of standoff distances shown.

Particle Filter for Localization

With the cosine model providing a directional estimate and uncertainty propagation providing the associated directional uncertainty, we are now in a position to address our original goal:

localizing an unknown neutron source. Several reported methods on localizing a radioactive source require a priori knowledge of the environment or large sensor networks, which are very limiting prerequisites for a field-deployable system. Our goal is to design a system which efficiently uses radiological measurements to inform the search algorithm and requires minimal *a priori* information. The simplest case we'll consider, in terms of motion planning, is scanning through a gas centrifuge uranium enrichment plant. Such a measurement campaign likely has a set raster pattern through the rows of centrifuges, where only the measurement time, spacing, and number of dwell points is variable.[10] For verifying the absence of radiation when no source is detected, we require broad coverage of the search environment. When a source is present, the robot must detect and converge toward the source. In this sense, in addition to a classical explore-exploit tradeoff, the problem may also be framed as task-switching, selecting an appropriate motion and measurement plan based on the observed situation. Working toward fully autonomous search capabilities, here, we focus primarily on detecting and converging toward the source.

While the Inspector Bot could directly follow the result of the cosine model, this does not effectively utilize past information. To incorporate past information along with the measurement uncertainty, we utilize a *particle filter*. Each "particle" represents a hypothesized source intensity and location and is assigned a weight according to a cumulative likelihood function. In the case of particle filtering *without* re-sampling,[11] the posterior $bel(x_t) = P(x|z)$ is estimated by updating the particle weights, w, according to Eq. 2, with measurements z and source parameters x.

$$w_t = P(z_t|x_t) \times w_{t-1} \tag{2}$$

The particle filtering method described requires a measurement model, $P(z|x)$, in order to assign weights to the particles, where the weight represents the probability of the corresponding particle, i.e., source hypothesis, yielding the acquired measurements. Assuming independence of measurements and source parameters, the joint probability of acquiring multiple measurements is given as the joint probability $\prod_i P(z_i|x_i)$. However, as the measurement uncertainty decreases or the number of measurements increases, the joint probability may tend toward zero. To avoid numerical errors, it is advantageous to convert to logarithmic likelihoods. For the probability mass function (PMF, discrete) or probability density function (PDF, continuous) given by $P(z|x)$, let $f(z|x)$ be the corresponding logarithmic likelihood function. In addition to numerical accuracy, the conversion to logarithmic likelihood functions is also convenient since the logarithm of a product is simply the sum of logarithms. This results in a great simplification of the distributions, since the logarithmic likelihood of several measurements can simply be added together. Further simplification is also possible when considering the normalization of the particle weights; any terms in the logarithmic likelihood functions which depend only on the measurement will be the same for all particles, and can thus be neglected in calculations since such constant terms will otherwise cancel-out.[12]

The PMF for the detected gross counts is defined in Eq. 3 as the Poisson distribution for detecting the counts z given the hypothesized source counts x. As a general characteristic

of Poisson distributions, the variance and mean are equal to the counts, so $\sigma_z^2 = z$. The corresponding logarithmic likelihood function is given by Eq. 4, which is further simplified, denoted by the primed function, to neglect terms which depend only on the measurement.

$$P(z|x) = e^{-x}\frac{x^z}{z!} \tag{3}$$

$$f(z|x) = -\ln(z!) + z\ln(x) - x \Longrightarrow f'(z|x) = z\ln(x) - x \tag{4}$$

The PDF for the estimated direction toward the source is defined in Eq. 5 as the von Mises distribution for detecting the direction z given the hypothesized source location x. In using the von Mises distribution, we assume that the concentration may be taken as the inverse of the variance, $\kappa \approx \frac{1}{\sigma_z^2}$. This approximation is valid for $\kappa \gg 0$ or, analogously, small σ_z^2. When this condition holds, the von Mises distribution may be rewritten, as shown in Eq. 5. Although the von Mises distribution for small κ more closely resembles a uniform distribution, for the sake of estimation, we assume the inverse relation between concentration and the standard flattened dispersion holds. The corresponding logarithmic likelihood function is given by Eq. 6, which is again simplified to neglect terms which depend only on the measurement.

$$P(z|x,\kappa) = \frac{e^{\kappa\cos(z-x)}}{2\pi I_0(\kappa)} \approx \frac{e^{\frac{\cos(z-x)}{\sigma_z^2}}}{2\pi I_0\left(\frac{1}{\sigma_z^2}\right)} \tag{5}$$

$$f(z|x) \approx \frac{\cos(z-x)}{\sigma_z^2} - \ln\left(2\pi I_0\left(\frac{1}{\sigma_z^2}\right)\right) \Longrightarrow f'(z|x) = \frac{\cos(z-x)}{\sigma_z^2} \tag{6}$$

For localizing and estimating the strength of an unknown neutron source, we utilize a prior distribution uniform in space and intensity. The weights of the particles are updated according to the distributions defined in Eq. 4 and 6. The result of particle filter described is visualized in Figure 5. In the simulated scenario shown, an unknown neutron source is located in open space. The Inspector Bot localizes the source over successive iterations, where each iteration consists of acquiring a new measurement, updating the particle weights, and moving a fixed step size toward the center-of-mass of the weighted particles. By leveraging the directional sensitivity of our Inspector Bot, we correctly estimate the location and intensity of a source with fewer measurements, as compared to particle filtering based only on gross counts. Our efforts in developing this particle filtering method into a more robust, deployable algorithm are ongoing. Two major thrusts are more advanced directional sensitivity with contextual information and motion planning. The former will be addressed by the six-detector system in conjunction with on-board sensors such as LIDAR. For the latter, we are considering scenarios with an initially undetectable source in which the Inspector Bot must begin by following a prescribed path. For instance, the Inspector Bot may follow a boustrophedonic raster pattern until the measurement uncertainty is sufficiently low for the particle filter to yield a reliable source estimate.

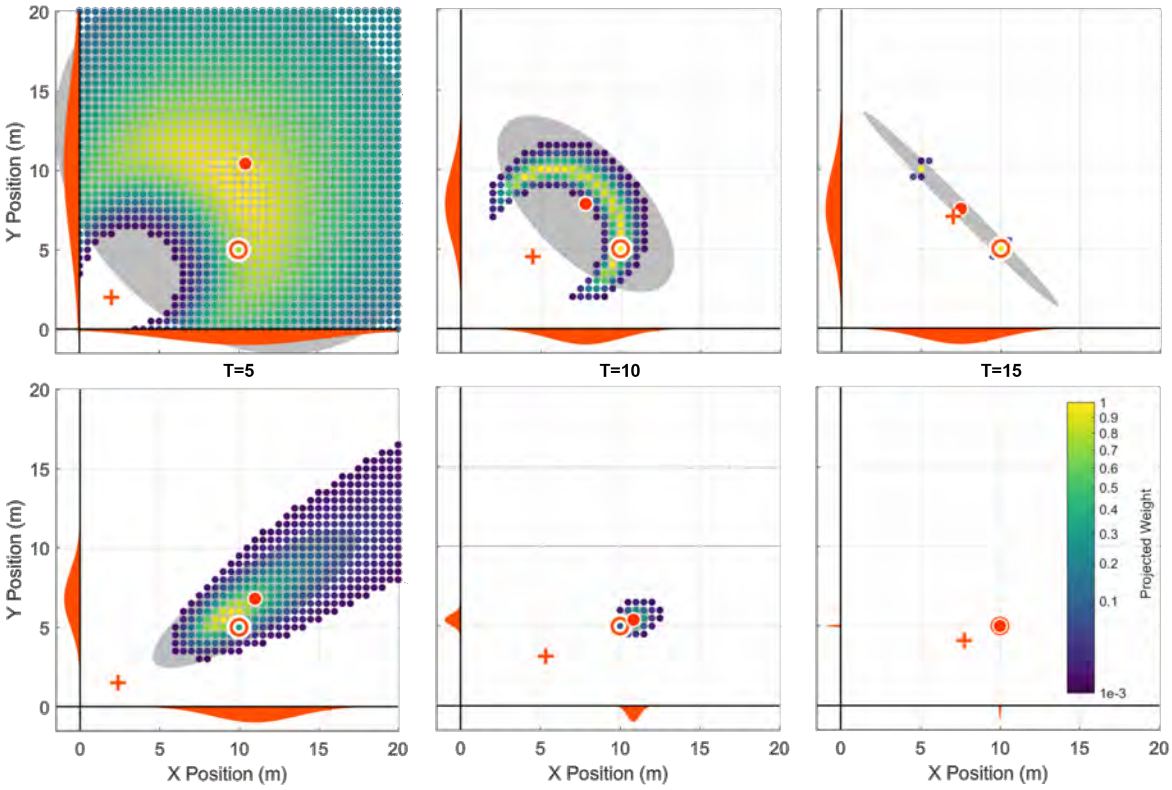

Figure 5: Snapshots of the particle filter for localization of an unknown source in a 400 m^2 search area. At each iteration, the robot (orange crosshair) takes a new set of measurements (i.e., the counts in each of the three detectors) and moves one meter toward the estimated source location (orange solid dot). The true source location is indicated by the orange ring at (10,5). Particles are initially uniformly distributed in space and intensity. Particles are colored by their respective weights, and the 95% confidence interval is highlighted in gray. **Above:** Particles are assigned weights according to the total counts. **Below:** Particle weights are computed as a function of both the total counts and the estimated direction. From left to right, snapshots of the weighted particles are shown for increasing iteration count. Evidently, with the directional sensitivity of our Inspector Bot, the true source is correctly identified in fewer iterations.

Discussion

Autonomous mobile robots, if designed properly, may be more effective and efficient and less intrusive than their human counterparts. We believe that our Inspector Bot, with the detection system comprised of boron-coated straw detectors encased in moderating high-density polyethylene, has significant potential to find application in nuclear safeguards and arms-control. Through a series of experimental measurements and MCNP modeling, we have examined the system performance of our Inspector Bot and characterized the directional sensitivity of the detection system. We also demonstrated how the directional sensitivity can be leveraged in the context of a particle filter for localizing an unknown neutron source. Future work on a six-detector version of the Inspector Bot will improve upon the directional sensitivity of the detector system by allowing us to more efficiently acquire data which may reveal

asymmetry and distortion in the detected counts caused by reflected neutrons. With further study, it may be possible to more accurately estimate the direction toward a source even with large nearby neutron reflectors. Further algorithm development and motion planning, additional experiments, and continued modeling will help achieve our vision of transitioning our Inspector Bot from a prototype to a deployable, autonomous inspector.

Endnotes

[1]Similar motivation was provided by the International Atomic Energy Agency, "Anonymized Report IAEA Robotics Challenge." 2018.

[2]More in depth discussion 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.

[3]Although our Inspector Bot currently has three BCS detectors, the HDPE moderator was designed with three additional holes (filled with HDPE plugs in the current configuration) for detectors to be added in the future. In principle, size is the only factor limiting the number of detectors.

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

[5]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 average of the two back detectors. In principle, if the position of a source is well-characterized, the front-to-back ratio may be further indicative of the energy of the source. We defer further investigation of this spectral sensitivity for future work.

[6]MCNP6.2 Release Notes, LA-UR-18-20808, Los Alamos National Laboratory, New Mexico, 2018.

[7]We thank Kyle Haske (Argonne National Laboratory) and Andrey Mozhayev (Pacific Northwest National Laboratory) for providing information on the source capsule configuration.

[8]For 5 m standoff distance, the bare source front-to-back ratio was 2.48 experimentally and 2.51 in MCNP. The moderated case was 3.43 experimentally and 3.32 in MCNP. Despite the agreement by this metric, the scaling of the absolute counts, which was calibrated to the bare source with 0.8% error, results in 13% error for the moderated case, suggesting further investigation will be required to fully understand the moderated source.

[9]With three detectors, there is systematic error in the estimated direction up to 5° (assuming zero background) when the "front" detector is rotated 30° away from the true direction. However, with six detectors, this intrinsic error is effectively eliminated as a result of improved curve-fitting.

[10]For instance, in "A Successive-Elimination Approach to Adaptive Robotic Source Seeking," IEEE Trans. Robot., 2020, Rolf et. al. proposed a method by which a robot repeats a fixed path and adjusts speed in real-time to ensure full coverage of the environment while maximizing information.

[11]Common practice is to 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.

[12]Both of the distributions used, Poisson and von Mises, may be approximated by a Gaussian PDF for certain conditions, i.e., when the detected counts increase and the variance in the directional estimate diminishes, respectively.