Estimating uranium hexafluoride cylinder contents using autonomous measurement systems

Mark S. Bandstra, Thomas D. MacDonald, Karthika Balan, Pei Yao Li, Kushant Patel, Emil Rofors, Marco Salathe, and Brian J. Quiter

Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA 94720, USA

June 2024

Abstract

In materials accountancy, radiological measurements of the nuclear material content of storage cylinders can be leveraged as a robust method of verifying declarations. These measurements are usually time-consuming and can result in appreciable radiation dose to personnel. In the near future it will be possible to task robotic systems with making such measurements. Such systems will use computer vision to orient themselves and locate the cylinders, and take measurements using onboard gamma-ray detectors. However, these systems will need real-time and near-real-time methods to assess whether the collected measurements are sufficient to determine the cylinder contents to a high degree of certainty, or whether more measurement locations and/or longer dwell times are necessary. In this work, we developed a simple physical model of a detector making a series of measurements around an arbitrary — but typically cluttered — arrangement of uranium hexafluoride cylinders and evaluated different algorithms for determining whether the containers were filled or empty, focusing on maximum likelihood (ML) and genetic algorithm approaches. We find that such approaches are promising, although there are degeneracies due to different possible uranium hexafluoride fill morphologies that significantly complicate the analysis.

1 Introduction

The measurement of storage containers in facilities to determine compliance with declarations is an important and well-studied area within nuclear safeguards [1, 2, 3]. Many measurements are currently performed manually *in situ*, with some automated systems also being used that require the containers to be moved into or along a rig [4]. With the advent of the latest robotics technology, there is now the opportunity to take advantage of autonomous, self-propelled systems to make these measurements. Such systems will be able to automate measurements in order to reduce the dose to personnel and free them to work simultaneously on other accountancy tasks.

We consider utilizing freely moving sensing systems, such as those carried by quadruped robots like the Boston Dynamics Spot [5]. Quadruped robots can traverse stairs, take elevators, and navigate various environments, making them good candidates for moving around stored material facilities. We assume the system is equipped with the sensors needed to make a 3-D map of the facility, including the locations and orientations of the containers of interest, and including the path and orientation of the system in the 3-D scene. New techniques are needed to analyze the data from such freely moving systems to take full advantage of

the 3-D scene information available to quickly verify the radiological signatures are what is expected, as well as to determine if more data is needed and from which locations in the scene.

In this work, we assume the following concept of operations. A map of a facility is provided to the inspector, along with the location of various storage containers and a manifest describing their contents. A robot-borne measurement system is brought to the facility and allowed to move throughout the facility. As it does so, the system builds its own 3-D map of the facility using, e.g., onboard lidar and cameras; identifies objects of interest; and collects radiation measurements using an onboard gamma-ray detector. Periodically during the inspection, the system will check what it has found so far against the manifest it was provided and present that information to the inspector. The system takes advantage of its 3-D mapping capabilities, its determination of object locations, and radiation measurements to estimate whether or not the radiation data it has collected is consistent with the manifest.

To make the problem more concrete, in this work we limit ourselves to the specific problem of measuring uranium hexafluoride (UF_6) cylinders in storage. This problem has been studied by many [6, 7, 8, 9, 10, 11, 12], and it poses unique challenges compared to other stored radioactive materials. UF_6 is a volatile substance, and some portion of it may be found distributed evenly throughout the inside surface of the cylinder via sublimation, while some may have settled into the bottom of the cylinder. When a large amount of the material has sublimated around the interior surface of a UF_6 cylinder, only limited information can be obtained about the total quantity [7, 13, 12, 14], which can make it difficult to fully assess the internal contents of the cylinder. This problem is further exacerbated if that thickness exceeds the skin depth of the photon of interest.

Although UF_6 measurements are often focused on measuring enrichment, we will not cover that topic here, assuming that information will be provided in a manifest. We focus instead on being able to measure *whether* the cylinder contains radioactive material, if so *how much*, and whether the amounts are consistent with the manifest. We will present and discuss two different methods for analyzing the data, one of which can be quickly calculated but is less accurate, and the other which is more accurate but slow, using a simple model of a measurement scenario to generate data.

2 Methods

Here we will present our simplified method for calculating the response a simulated detector would have to radioactivity in a set of hypothetical models. We then use these response calculations as inputs to two analysis approaches we examined, which differ in precision and computational requirements.

2.1 Calculating responses to material in cylinders

The basis of the methods presented here is a 2.5-D model meant to estimate the measurement of UF_6 cylinders. We assumed the facility contains cylinders of type 30B, they are stored horizontally (central axis parallel to the ground), and the uranium is enriched to 4.95%. The radiation detector only moves at a fixed height relative to the UF_6 cylinders, i.e., in the plane perpendicular to its axis that would slice the cylinder, hence the 2-D nature of the model (the other half dimension comes from using some z dimension information later). The detector will also be a simple, non-imaging detector — we assumed a perfectly isotropic response of 37.2 cm^2, which is approximately the effective area of a 3 $\times$ 3-inch NaI(Tl) cylindrical detector for the ^{234m}Pa line at 1001 keV. We will only consider the 1001 keV line since it is the most penetrating and thus sensitive to the bulk of the material, but one could also analyze the 186 keV line from ^{235}U to examine enrichment as well.

UF_6 can undergo various phase changes during storage due to movement of the cylinder and thermal cycling, leading to complex deposition patterns inside the cylinder [7, 12, 14]. To quantify the morphology

of the UF_6 inside each cylinder, we used the fill profile quantity x defined by Berndt et al. [7] to parameterize the fill shape. The value of x gives the fraction of the material deposited around the inner wall of the cylinder in a concentric ring, while the other portion $(1 - x)$ lies in a circular segment at the bottom of the ring.

The data used in the model were obtained as follows. Attenuation lengths for 1001 keV photons were obtained from XCOM [15], material properties where obtained from ref. [16], and branching ratios and half-lives from ref. [17]. The dimensions of 30B cylinders are given in ref. [18] and summarized in ref. [12]. Specifically, we used an outer diameter of 76.2 cm, a length of 193 cm, and a 1/2-inch steel outer wall thickness.

2.1.1 Calculating 2D cylinder fill profile

Letting f be the fraction of the total cylinder volume filled by UF_6 (total circle area in 2D), and x to be the fill profile, then we can calculate the physical parameters of the 2D fill profile. Figure 1 shows the geometry and definitions used to find the thickness of the outer concentric ring due to the sublimation (t), the angle subtending the circular segment that contains the material settled in the bottom (θ), and the depth of the material in the circular segment inside the outer ring (d).

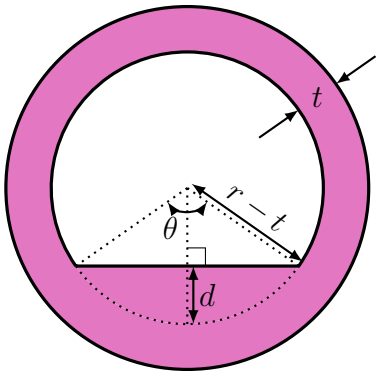

$$t = r\left(1 - \sqrt{1 - xf}\right)$$

$$\theta - \sin\theta = 2\pi\frac{(1-x)f}{(1-xf)}$$

$$d = r\sqrt{1 - xf}\,[1 - \cos(\theta/2)]$$

Figure 1: Diagram explaining the method of calculating the 2D profile shapes of UF_6 in a cylinder given the fill fraction f and the fill profile x.

2.1.2 2-D ray casting

Now with the (UF_6) contents of a cylinder described by f and x, we need to be able to estimate the photon flux due to the cylinder material at any point in space. To do this, we first take a sample of points inside the circle and discard any that are in the air cavity. The sampling we chose was the sunflower arrangement [19] so that all of the points would be approximately the same distance from their nearest neighbors and to minimize sampling artifacts we observed using polar grids. Figure 2 shows an example of such a sampling, though for 500 points for illustrative purposes, rather than the 2,000 points that were used in the analyses herein.

Additionally, we needed to take each ray from a point inside the UF_6 to the detector point and determine how much of which materials it passed through. This is the process of ray casting. To accomplish this, we wrote routines in python to find intersections with and lengths traversed through annuli and circular segments. If μ is the attenuation coefficient and r the distance traversed through each material (indicated as subscripts), the flux from a point at (x, y) to a detector at $(0, 0)$ is

$$\text{flux}_{2D} \propto \frac{1}{4\pi(x^2 + y^2)} \exp\left(-\mu_{\text{UF}_6} r_{\text{UF}_6}\right) \exp\left(-\mu_{\text{steel}} r_{\text{steel}}\right) \exp\left(-\mu_{\text{air}} r_{\text{air}}\right). \quad (1)$$

Finally, to approximate the 3-D extent of the cylinders away from either side of the central plane ($z = 0$), we sampled points along the length of the cylinder in the z dimension and scaled the dimensions of the previous formula by the ratio of the 2-D distance to the 3-D distance:

$$\text{flux}_{3D} \propto \frac{1}{4\pi(x^2+y^2+z^2)} \exp\left(-\mu_{\text{UF}_6} r_{\text{UF}_6} \eta\right) \exp\left(-\mu_{\text{steel}} r_{\text{steel}} \eta\right) \exp\left(-\mu_{\text{air}} r_{\text{air}} \eta\right) \tag{2}$$

where $\eta = \frac{\sqrt{x^2+y^2+z^2}}{\sqrt{x^2+y^2}}$. We then averaged over the values of z that span the height of a container using 100 samples in that dimension.

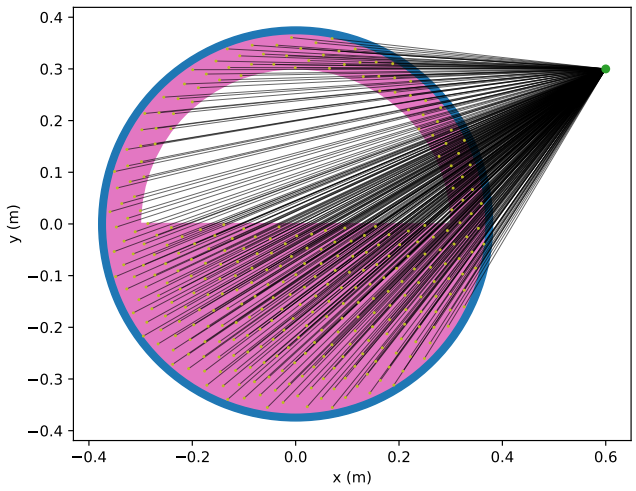

Figure 2: An example of the 2-D ray casting performed to calculate the response for a single cylinder for a single detector measurement. The UF_6 material is shown in magenta, and the steel shell is shown in blue. The detector measurement point is the upper right point where all of the lines converge, while the yellow points inside the filled sections of the cylinder are the sampled points.

2.2 MLEM estimation of cylinder contents

To rapidly estimate the radioactive contents of an arrangement of cylinders, we developed a simple linear model that we can solve quickly and that can also render preliminary estimates of the uncertainties of each cylinder's contents. For the model, we first choose some nominal cylinder configuration as the baseline; for example, a UF_6 cylinder that is two-thirds full with x=100% sublimation. Assume we have N cylinders in the scene and M measurement points. For each detector measurement point, we calculate the response to each cylinder according to the previous section to obtain the expected mean count rate in the chosen photopeak region per cylinder activity. We then place these N mean count rates into a row of an $M \times N$ response matrix we will call $\mathbf{A}_{\text{cyl}}$, and we continue the calculation for all M measurement points. Placing the measured gamma-ray counts in a vector $\mathbf{y}$, and assuming cylinder contents can be modeled by scaling relative to the nominal cylinder's activity, we get the following equation:

$$\mathbf{y} \approx \bar{\mathbf{y}} \equiv (\alpha_{\text{nom}} \mathbf{A}_{\text{cyl}} \boldsymbol{\theta}_{\text{cyl}} + \mathbf{b}) \odot \boldsymbol{\Delta t}, \tag{3}$$

where α_{nom} is the activity of the nominal container, $\boldsymbol{\theta}_{\text{cyl}}$ are the unknown coefficients that scale each of the cylinders' apparent activities relative to the nominal cylinder's activity, $\mathbf{b}$ are the mean background count rates per measurement, $\odot$ represents element-wise multiplication, and $\boldsymbol{\Delta t}$ are the integration times at each measurement point. So if all of the cylinders are the same as the nominal assumed cylinder, then all of the elements of $\boldsymbol{\theta}_{\text{cyl}}$ would be one.

To proceed further, we will assume the background rate is unknown and a constant (i.e., $\mathbf{b} = b\mathbf{1}_M$, where $\mathbf{1}_M$ is a column vector of M ones). We augment $\mathbf{A}$ and $\boldsymbol{\theta}$ to absorb b:

$$\mathbf{A} \leftarrow \left[\begin{array}{cc} \mathbf{1}_M & \alpha_{\text{nom}}\mathbf{A}_{\text{cyl}} \end{array} \right], \quad \boldsymbol{\theta} \leftarrow \left[\begin{array}{c} b \\ \boldsymbol{\theta}_{cyl} \end{array} \right]. \tag{4}$$

Since the system measures event counts, we expect the data to be distributed according to the Poisson distribution, with $\bar{\mathbf{y}}$ on the right-hand side of equation (3) being the mean counts. Therefore, we use the negative log likelihood as the cost function to minimize:

$$-\log L(\mathbf{y}|\boldsymbol{\theta}) = \sum_{i=0}^{M-1} \left(\bar{y}_i - y_i \log(\bar{y}_i) + \log(y_i!)\right). \tag{5}$$

To minimize this cost function we can use maximum likelihood expectation maximization (MLEM) [20, 21]. Starting with any non-negative vector $\hat{\boldsymbol{\theta}}^{(0)}$, MLEM gives the iterative update rule

$$\hat{\boldsymbol{\theta}}^{(k+1)} \equiv \hat{\boldsymbol{\theta}}^{(k)} \odot \left[\frac{\mathbf{A}^\top \left(\boldsymbol{\Delta t} \odot \frac{\mathbf{y}}{\mathbf{A}\hat{\boldsymbol{\theta}}^{(k)}} \right)}{\mathbf{A}^\top \boldsymbol{\Delta t}} \right]. \tag{6}$$

We iterated until the arbitrary stopping condition that all elements of $\boldsymbol{\theta}$ change by less than 10^{-6}.

Once we have the MLEM solution, we can also use the properties of the problem to estimate the statistical uncertainties of each element of $\hat{\boldsymbol{\theta}}$. We will do this by calculating the Fisher information matrix, $\mathbf{F}$, defined as the expectation value of the Hessian of the negative log likelihood function evaluated at $\hat{\boldsymbol{\theta}}$, which is approximately

$$\mathbf{F} \approx \mathbf{A}^\top \operatorname{diag}\left(\frac{\boldsymbol{\Delta t} \odot \boldsymbol{\Delta t}}{\mathbf{A}\hat{\boldsymbol{\theta}}} \right) \mathbf{A}. \tag{7}$$

Expecting that at least hundreds of counts will be measured, then the estimator should approach the Cramér-Rao lower bound (CRLB) [22], which means that the covariance of the estimator is approximately

$$\operatorname{var}(\hat{\boldsymbol{\theta}}) \approx \mathbf{F}^{-1}. \tag{8}$$

And the diagonal of this matrix provides estimates of the variances of each of the coefficients.

2.3 Estimation of cylinder contents using a genetic algorithm

For a more accurate but slower approach, we attempted to find the best values of both f and x for all cylinders that best fit the data. But because the cylinder contents determine the attenuation of photons from neighboring cylinders, the problem is not convex. To do this, we divided the possible f and x values into discrete sets — 0%, 33%, and 67%; and 0, 0.25, 0.50, 0.75, and 1.0, respectively. We used a genetic algorithm to efficiently search the space of possible solutions.

The genetic algorithm was implemented as follows. First, 50 random configurations of f and x where chosen. Each configuration is an "individual" in the "population". At each generation, the Poisson negative likelihood was calculated between the predicted counts of each individual and the actual measured counts from the simple model. The top 10% individuals were considered "elite" and left untouched; the next 20% in ranking were considered fit enough to keep; while the lowest 70% were discarded. The empty slots were filled by "crossing over" randomly chosen pair of individuals from the top 30% — i.e., their f and x values were randomly interchanged to create new individuals. Finally, the bottom 90% (non-elites) were mutated, where each value of f or x in the individual could be randomly changed to another discrete value with a probability of 20%. After 20 generations, the solutions were examined.

3 Results

To explore the implications of the simple model and the different approaches to solving for cylinder contents, we chose a simple scenario. In the scenario, the detector travels directly above a two-layer stack of cylinders as shown in Figure 3. We assumed an integration time of 0.25 s and the detector moving at a constant speed of 0.15 m/s.

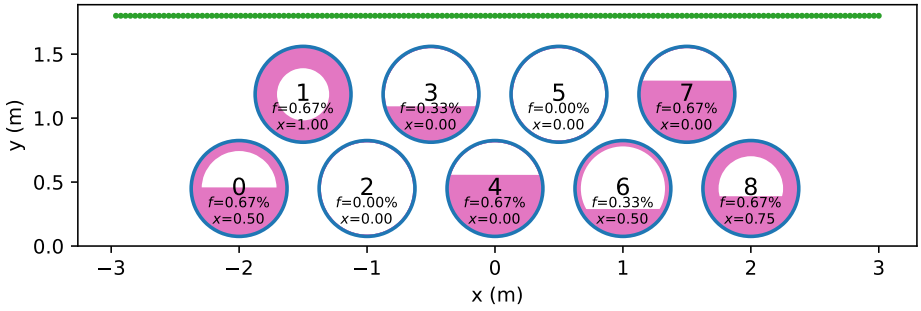

Figure 3: Scenario used in the main analysis. The detector moves in a line shown with green points, measuring the 1001 keV photons emitted by a set of cylinders, in some cases stacked on top of others.

The MLEM estimation from Section 2.2 was used, assuming a nominal cylinder type of $f = 0.67$ and $x = 1$ for all cylinders. The estimate was performed in approximately 4.5 seconds on a 2023 MacBook Pro with an M2 chip, most of which time was spent performing the ray casting needed for the response matrix. Figure 4 shows the "true" count rate calculated for the scenario, the Poisson sampling of that count rate, which simulates a measurement, the estimated count rate fit using the MLEM method, and contributions to the true and fit rates from each UF_6 cylinder. As can be seen, the fit is qualitatively close to the true value, but some differences remain due to the slight differences in emission and attenuation from the true distribution of cylinder contents versus the nominal assumption.

The estimated values of $\boldsymbol{\theta}$ and their statistical uncertainties are shown in Figure 5. The "true" values plotted alongside are the ratios between the maximum response for the actual configuration versus the nominal one. This plot reveals that the method correctly found that all of the cylinders contain UF_6 except 2 and 5, and that the activities contained in those are consistent with zero. Few of the estimates are statistically consistent with the nominal value of 1, but none of the cylinders are actually in the nominal configuration of $f = 0.67$, $x = 1.0$, so that is not unexpected.

The genetic algorithm, on the other hand, took approximately 100 minutes to run on the same computer, and the portion of the calculation that took the most time was again the calculation of the detector responses

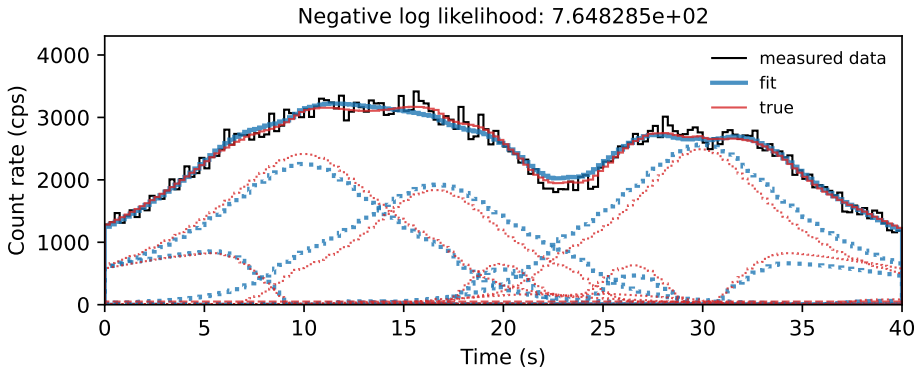

Figure 4: The MLEM solution assuming all cylinders are filled. The contributions to the true and estimated count rates from each of the seven cylinders individually are shown with dotted lines.

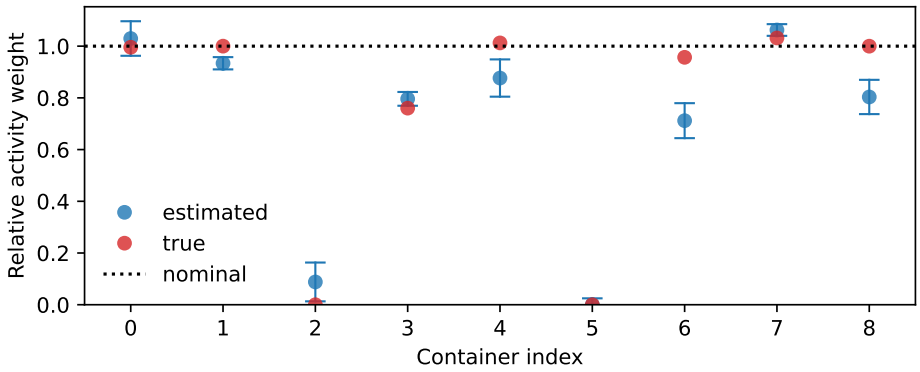

Figure 5: The MLEM coefficients $\hat{\boldsymbol{\theta}}$ for the cylinders in the scenario compared to the "true" values defined as the ratio of the maximum of the actual response for each cylinder to the maximum of the nominal response. Their estimated uncertainties shown are 1-sigma. A coefficient of one means that the emission is approximately that of the nominal cylinder configuration assumed.

for each new configuration of cylinder contents. The fitness scores of the population after 20 generations are shown in Figure 6. The Poisson negative log likelihood values were converted to equivalent Gaussian z-scores by assuming they were from a chi-squared distribution with $M - 1$ degrees of freedom (due to M bins but fitting one parameter, the constant background level). The z-scores reveal that the best few solutions in the population are at approximately 1-sigma, meaning the solutions are not able to be rejected at the 95% confidence level.

The top scoring solution is shown in Figure 7. This solution is close to the true cylinder configurations but not exactly the same. The actual count rates predicted by this solution are shown in the same figure, which reveals a significantly better agreement than the MLEM solution.

4 Discussion

The 2.5-D model used here is simple but intended to capture enough complexity to enable us to explore ways of approaching the problem of freely moving autonomous systems surveying a storage facility.

The MLEM method for rapidly estimating the contents of cylinders yielded results that are close to the

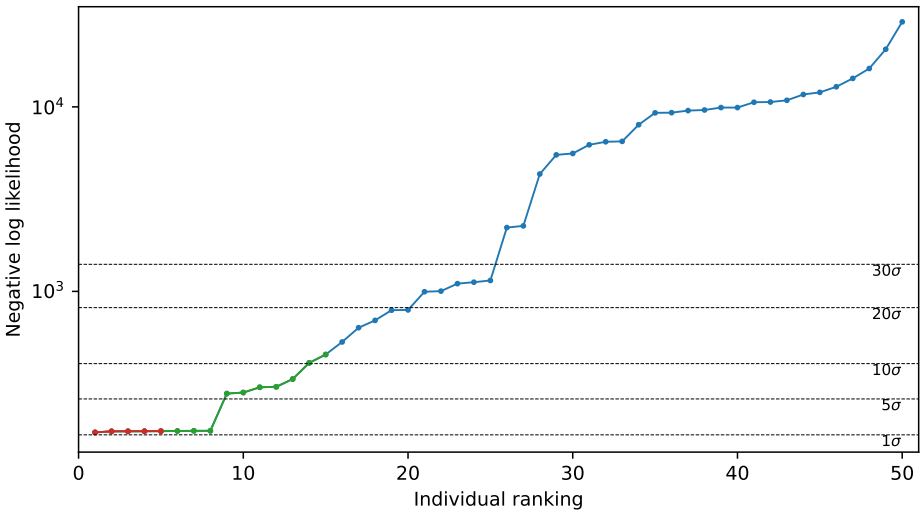

Figure 6: The ranking of fitness scores of all of the individuals in the genetic algorithm population after 20 generations. The red individuals would fall in the 'elite' category, while the green individuals would be retained and mutated in a subsequent generation.

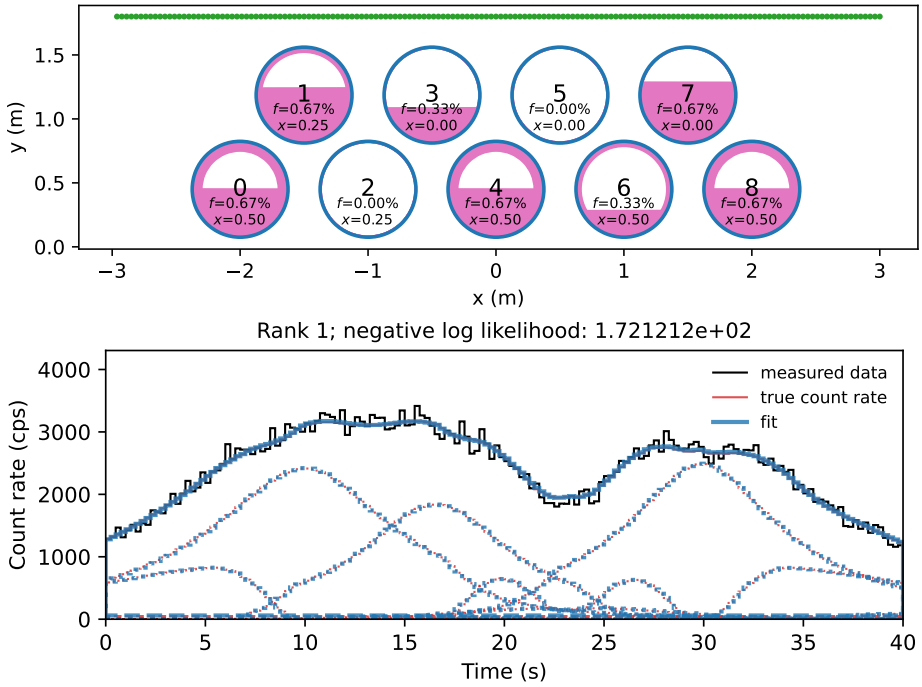

Figure 7: The best solution found by the genetic algorithm (top), and the true count rates compared to the count rates predicted by this solution (bottom).

true solution, and it requires only one expensive detector response calculation. However, because the method only considers the emission strength from each cylinder rather than their fill profiles, its accuracy is limited; however we expect to have *a priori* information about cylinder contents in the form of a facility's manifest

to mitigate this problem. The method also has the benefit of estimating the variance of the solution, which may be useful information as feedback for path planning, e.g., enabling the inspector (or the system itself) to quickly decide whether it needs a second pass, or for deriving useful heuristics.

A genetic algorithm was also presented that was able to obtain a closer fit to the data than the MLEM method but at the cost of increased computation time due to the large number of detector response calculations performed. However, even this algorithm could not correctly reproduce the original configuration because UF_6 cylinders with a thick enough layer of UF_6 around the inner wall have nearly identical emission and attenuation of 1001 keV photons. Since the mean free path in UF_6 is approximately 2.8 cm, any cylinder with a configuration that has an inner ring of material that is thicker than about a few times the mean free path can never be distinguished. This degeneracy is of course a well known problem in UF_6 measurements, especially for enrichment where the problem is even worse because the 186 keV photon from ^{235}U is vital to the estimate but has an even shorter mean free path [6, 12, 11].

5 Conclusion

We explored methods of using gamma-ray detector data collected by a freely moving autonomous system to reconstruct the contents of storage containers. With a simple 2.5-D model, we demonstrated how, even for a non-imaging gamma-ray detector, a simple MLEM approach could be useful to rapidly perform the reconstruction, while a slower genetic algorithm could perform the reconstruction more accurately. We challenged the approaches with a stack of cylinders with varying UF_6 configurations, though we expect actual cylinders in a facility to be similar to each other due to similar conditions, and thus some free parameters could be eliminated in practice.

The approaches presented here show the challenges associated with real time or near-real time assessments of stored radioactive material using a series of measurements from non-imaging radiation detector at stand-off distances that would be reasonable for an autonomous robotic system, in particular a trade-off between accuracy and computational speed. For the simple 2.5-D model used here, much of the computational cost was in CPU-bound python code to calculate the detector response matrix. There are likely speedups available through better-optimized CPU code and adapting the code to run on the GPU, which was not attempted here.

Another area of improvement is in moving to three dimensions. Obviously the 2.5-D model is highly simplified, and 3-D measurements could contain more information about the cylinders. Three-dimensional calculations would be even more complex but are beginning to be explored for this project. Existing open-source code for performing 3-D ray casting on GPUs should mitigate some of the computational costs.

Finally, additional information and complexity will come from the use of collimated or imaging detectors instead of a non-imaging detector, as was explored here.

Although UF_6 cylinders were explored, the goals of this project are to be able to rapidly assess a large number of possible stored materials. Near future work involves taking field measurements of depleted uranium cylinders and completing an analysis chain that starts with collecting multi-sensor data with a robot, to generating 3D models of the scene, to calculating detector response models and estimating using the methods explained here. Using this information for autonomous path planning is also being pursued.

6 Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory (LBNL) under Contract DE-AC02-05CH11231. The project was funded by the U.S. Department of Energy, National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development. We also thank Jayson Vavrek for his helpful comments.

References

[1] Nuclear Material Accounting Handbook. Technical report, IAEA, 2008.

[2] Use of Nuclear Material Accounting and Control for Nuclear Security Purposes at Facilities. Technical report, IAEA, 2015.

[3] International Safeguards in the Design of Uranium Conversion Plants. Technical report, IAEA, 2017.

[4] L. E. Smith et al. An unattended verification station for UF6 cylinders: Field trial findings. *NIM A*, 874:127–136, December 2017. ISSN 0168-9002. doi: 10.1016/j.nima.2017.08.003.

[5] Boston Dynamics. Spot - The Agile Mobile Robot, 2024. `https://bostondynamics.com/products/spot/`, Last accessed on 2024-06-18.

[6] H. A. Smith. The measurement of uranium enrichment. Technical Report NUREG/CR-5550, Los Alamos National Laboratory, 1991.

[7] R. Berndt, E. Franke, and P. Mortreau. 235U enrichment or UF6 mass determination on UF6 cylinders with non-destructive analysis methods. *NIM A*, 612(2):309–319, 2010.

[8] E. K. Mace and L. E. Smith. Automated nondestructive assay of UF6 cylinders: Detector characterization and initial measurements. *NIM A*, 652(1):62–65, 2011.

[9] K. A. Miller. Conceptual Ideas for New Nondestructive UF6 Cylinder Assay Techniques. Technical Report LA-UR-12-21067, Los Alamos National Laboratory, 2012.

[10] B. Canion et al. Standoff enrichment analysis of UF6 cylinders. *NIM A*, 954:161342, 2020.

[11] A. T. Greaney et al. Comparison of gamma-ray spectral analysis methods for thick-walled UF6 cylinders. *NIM A*, 977:164291, 2020.

[12] N. McFerran et al. Gamma-ray spectrum variations for surface measurements of uranium hexafluoride cylinders. *NIM A*, 961:163675, 2020.

[13] A. T. Greaney et al. Heterogeneity effects on nondestructive assay measurements of enrichment in UF6 cylinders. *ESARDA Bulletin*, 61(1), 2020.

[14] A. T. Greaney et al. Thermal impacts on nondestructive analysis measurements of uranium hexafluoride. *Journal of Radioanalytical and Nuclear Chemistry*, 330(1):357–365, 2021.

[15] M.J. Berger et al. XCOM: Photon cross sections database. Technical report, NIST Standard Reference Database 8, 2010.

[16] R. S. Detwiler et al. Compendium of Material Composition Data for Radiation Transport Modeling. Technical Report PNNL-15870-Revision-2, 2021.

[17] National Nuclear Data Center, 2024. URL `https://www.nndc.bnl.gov/nudat/`.

[18] ANSI N14.1-2023 - Nuclear Materials - Uranium Hexafluoride - Packagings for Transport.

[19] H. Vogel. A better way to construct the sunflower head. *Mathematical Biosciences*, 44(3):179–189, 1979.

[20] L. A. Shepp and Y. Vardi. Maximum Likelihood Reconstruction for Emission Tomography. *IEEE Trans. on Medical Imaging*, 1(2):113–122, 1982.

[21] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *Journal of Computer Assisted Tomography*, 8(2):306–316, 1984.

[22] A. Stuart, K. Ord, and S. Arnold. *Kendall's Advanced Theory of Statistics, Sixth Edition*, volume 2A: Classical Inference and the Linear Model. Wiley, 2004.