Scene-informed optimal measurement positions for quantitative safeguards measurements

K. Knecht, M. S. Bandstra, G. Batie, J. Daughhetee, M. Folsom, D. Gunter, D. Hellfeld, T. H. Y. Joshi, M. Salathe, K. Schmitt, K.-P. Ziock and K. Vetter

Abstract

Portable radiation detection systems can be equipped with contextual sensors to allow free-moving 3D gamma-ray imaging through scene data fusion (SDF). While developed for free-moving measurements, SDF also has applications in cases that require multiple static measurements to generate quantitative 3D images, where the scene information can improve results. The scene information captured by these devices can be leveraged to determine optimal measurement positions for quantification measurements, such as in safeguards applications, where limited time is provided to quantify nuclear material present.

In this work the scene information obtained from contextual sensors is used to construct a 3D voxel grid. A 3D object of interest within this grid is selected by an operator either manually or automatically through an object detection machine learning framework. A number of potential measurement positions and orientations (poses) around this object of interest are identified, and the sensitivity to each of the object voxels from each of these poses is computed. Ray tracing is used to estimate any attenuation term, and the detector response for the Compton imaging modality is used in computing sensitivity. The combination of positions which results in maximum total sensitivity with minimum variance in sensitivity across object voxels are identified as the optimal measurement positions. For large measurement spaces the optimization is computationally intensive, so a genetic algorithm is used to accelerate the computation of optimal positions.

We first demonstrate the feasibility of the optimization framework with a toy model. We then use a SDF-enabled gamma-ray imager to find optimal measurement positions for a mock material holdup loop with distributed sources. We present our approach and results by performing an optimization to produce quantitative Compton images in these low-count-rate environments and the minimization of associated uncertainties.

I. Introduction

THE detection, characterization, localization, and mapping of gamma-ray source distributions in real-world environments is important for many applications including safeguards, emergency response, and nuclear contamination remediation. Gamma-ray imaging presents a passive, non-destructive method for providing a quantitative measure of nuclear materials present.

Portable detection systems can be equipped with contextual sensors to allow for free-moving 3D gamma-ray imaging in a method called scene data fusion (SDF) [1]. SDF presents a non-destructive means of analyzing nuclear material present in a variety of environments. SDF captures relevant scene data including a 3D map of an area, which can be used to inform operator decisions and enhance quantitative imaging results. The scene information can be leveraged to determine optimal measurement positions to reduce measurement time while still capturing enough data to quantify to desired statistical uncertainties. This is a beneficial technology for safeguards inspections, particularly when a limited amount of time is available to perform measurements, ideally with minimal disruption to facility operation.

This material is based upon work supported by the Department of Energy National Nuclear Security Administration through the Nuclear Science and Security Consortium under Award Number(s) DE-NA0003180. This support does not constitute an expressed or implied endorsement on the part of the United States Government.

K. Knecht, G. Batie, and K. Vetter are with the Nuclear Engineering Department at the University of California, Berkeley, Berkeley, CA 94720 USA (e-mail: kalie@berkeley.edu).

M. S. Bandstra, M. Folsom, D. Hellfeld, T. Joshi, M. Salathe and K. Vetter are with the Nuclear Science Division at Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA.

J. Daughhetee, K. Schmitt, and K.-P. Ziock are with the Physics Division at Oak Ridge National Laboratory, Oak Ridge, TN 37830 USA.

D. Gunter is with Gunter Physics, Inc. at 6221 Trinity Dr., Lisle, IL 60532 USA

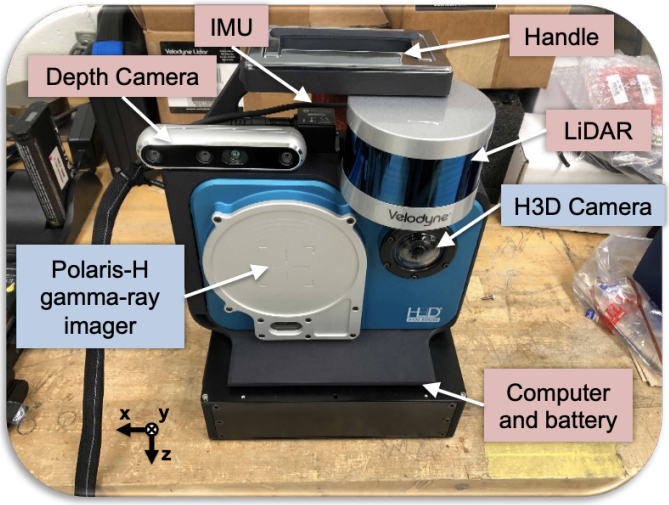

Fig. 1 The H420-LAMP imaging system and its coordinate system. The sensors labeled in blue are included in the commercially available H420. The sensors labeled in red are part of the LAMP package.

II. Methods

A. *H420-Lamp*

The system used to demonstrate this work is the H420-LAMP imaging system, shown in Fig. 1. This device consists of the commercially available H420 quad CdZnTe (CZT) detector, designed by H3D as a 2D static imager capable of coded-aperture and Compton imaging [2]. The imager has been retrofitted with the Localization and Mapping Platform (LAMP) developed at Lawrence Berkeley National Laboratory, which includes a RealSense depth camera, inertial measurement unit (IMU), and Velodyne LiDAR [3]. The device also includes and an onboard computer to perform Simultaneous Localization and Mapping (SLAM) with Google Cartographer [4], enabling the system to estimate its position and orientation and generate a 3D map of its surroundings. For generating the 3D map, the system is hand-carried through the scene. For point and dwell static measurements, the imager is placed on a tripod.

B. *Scene Data Fusion*

SDF has been extensively demonstrated in source search and non-quantitative mapping scenarios [1] [5] [6] [7]. Recently methods have been developed to quantify these measurements in point and distributed source environments [8]. These tools present a way to evaluate the amount of material present, which is needed for safeguards applications. The scene information gathered by the system can be further leveraged to enhance imaging capabilities, especially in cases requiring assessment of weak radiological signatures.

Previously SDF has been used solely for free-moving 3D imaging, but this work seeks to enhance the use of the contextual information for static 3D imaging. In this work, the H420-LAMP is walked around the scene where the LAMP package captures and reconstructs scene data to create point clouds, which describe the objects in the scene as points in the 3D imaging space. The imaging space is divided into a voxelized grid, where a voxel is considered to be occupied if the number of point cloud points which reside in it is greater than some cutoff value. The LiDAR point cloud only renders surfaces of objects, so voxels interior to objects are assumed to be empty unless manually set as occupied. The voxel grid is then used to optimize static measurement positions of the imager for quantifying activity emitted from a radioactive source. The occupancy information provided by the voxel grid is used to determine the position of the volume of interest, as well as feasible measurement positions that have line of sight to the volume. If material types are known, then the occupancy information of the voxel grid can also be used to estimate attenuation of intervening material to properly quantify material present. Some work has already been done to estimate attenuation of intervening materials in SDF applications [9].

C. Identification of volume of interest

This work assumes that the object being imaged occupies some volume within the voxel grid, which can consist of one or more occupied voxels. Therefore, the voxels which contain the imaged object must be identified. This is currently done manually, by indentifying the x-y-z limits of the object in the context of the 3D map generated in the scene survey. This region of the global voxel grid is cropped and identified as the object voxel grid.

D. Viable Measurement Position Selection

Once an object of interest is selected, a grid of potential measurement locations is automatically constructed around the object voxel grid between a user-provided maximum standoff distance and a minimum standoff distance such that all object voxels are in the coded aperture field of view of 86°x86°. We include this constraint because the imager is capable of coded aperture imaging, but the coded aperture imaging modality is not used in this work.

Some measurement positions are nonphysical or undesirable for a counting measurement. To remove these points, the measurement positions which reside in occupied voxels are removed. A ray-tracing algorithm based on [10] (pseudo-code in [9]) is used to determine if at least one object voxel is visible at each potential measurement position (some object voxels may be occluded by the object itself), and positions where no object voxels are visible are removed.

Once the potential measurement x-y-z positions are identified, a discretization in orientation around the z-axis of the imaging system (coordinate system shown in Fig. 1) may be defined. Orientations around the x- and y-axis are not included due to physical constraints of placing the imager on the tripod. Each potential orientation is considered for each of the x-y-z positions, but generally those which point toward the object of interest are desirable. Each combination of x-y-z positions and orientations is referred to as a potential 'pose' for the imager. Each of the potential poses are evaluated to serve as the the optimal imager placement for static imaging to quantify the volume of interest.

To produce a 3D image, the imager must be placed at more than one measurement position. Therefore, the desired outcome is an optimal combination of measurement poses. A permutation is performed to find all possible combinations of measurement poses. Each pose may be represented multiple times in the combination, representing longer dwell times at that location.

E. Measurement Pose Optimization

This approach seeks to find the combination of measurement poses which have maximum sensitivity to the volume of interest while having uniform sensitivity to each of the object voxels. Therefore, this is a multi-objective optimization that seeks to maximize the average sensitivity to object voxels while minimizing the variance in sensitivity of the object voxels across the combination of measurement poses.

1) Sensitivity: The sensitivity to each object voxel at each potential measurement pose is computed as:

$$S = R \times \frac{1}{r^2} \times e^{-\mu_{unocc} r_{unocc}} \times e^{-\mu_{occ} r_{occ}}$$

where:

- R is the expected detector response for double gamma-ray events at the desired energy at the solid angle between the given detector pose and an object voxel,
- r is the distance between the center of the object voxel and the detector position,
- μ_{unocc} is the linear attenuation coefficient of the provided gamma-ray energy in air at STP,
- r_{unocc} is the path length of air between the object voxel center and the measurement position,
- μ_{occ} is the linear attenuation coefficient of the provided gamma-ray energy in occupied voxels, and
- r_{occ} is the path length of occupied voxels between the object voxel center and the measurement position. Note that $r_{object} + r_{air} = r$.

The detector response is generated through experimentally benchmarked Geant4 simulations [8]. The polar and azimuthal angles between the measurement pose and source voxel are computed and the response is interpolated from response data for the desired gamma-ray energy. The ray-tracing algorithm previously described is used to find the occupied and unoccupied path lengths. The occupied attenuation coefficient is manually provided for intervening materials.

2) Brute-force method for finding best combination of poses: The brute-force method requires the optimization to be performed on all possible combinations of measurement poses. In this method, the sensitivities to each voxel at each potential measurement pose are first computed. The sensitivities to each voxel across each measurement pose in a combination are then summed to obtain a total sensitivity to that voxel in that combination of poses. The mean and variance of the summed object voxel sensitivities are then computed, which are the metrics used for the multi-objective optimization.

In multi-objective optimization, the optimization of one metric may be at the expense of another. To find the candidate optimal solutions we use the Pareto front, a concept borrowed from economic and game theory which seeks to find solutions that 'help' one metric without 'hurting' the other [11]. The Pareto front represents several combinations of measurement poses which may considered to be optimal solutions to this problem. There are a variety of methods to choose the desired candidate from the Pareto front, including the Nash bargaining solution [12] and the Kalai-Smorodinsky bargaining solution [13]. In this analysis, the Nash bargaining solution is used to select an optimal combination of measurement poses from the Pareto front to place the imager.

3) Algorithmic methods for finding best combination of poses: In the case of many potential measurement poses, the brute-force approach results in a large computational burden (e.g. for 1000 potential measurement poses and 3 optimal poses, there are over 1.67×10^8 possible combinations of poses which must be computed). In these cases, it is desired to use an optimization method such as a genetic algorithm (GA) to accelerate these calculations to facilitate real-time operation.

In a GA problem, a population of individuals are iteratively evaluated by a fitness function, where each iteration is called a 'generation'. The initial population is randomly selected from a global population. The 'fit' individuals in each generation are selected, where their 'genes' may be modified through crossover between different individuals (controlled by the crossover parameter) and random insertion from the global population (controlled by the mutation parameter) to form the next generation. The iteration continues until a maximum number of generations has been reached or the population converges to a satisfactory fitness level.

The Non-dominated Sorting Genetic Algorithm (NSGA-II) [14] from the PyGMO python package [15] is used to accelerate optimization. To formulate the optimization as a GA problem, each measurement pose is treated as a 'gene' and each combination of measurement poses is treated as an 'individual'. The global 'population' of individuals consists of all possible combinations of measurement poses. Only a small fraction of the global population is present in an initial population randomly sampled by the GA. The population is sized such that most of the potential poses are represented in the sample at least once, and the mutation parameter is set sufficiently high to increase the likelihood that those poses not initially represented are mutated into the population at some point in the optimization. The crossover parameter is left at the default value. For each generation of the GA, the metrics of interest are computed on-the-fly for each of the combinations of poses in that generation. The 'fittest' combinations are passed to the next generation, where individual measurement poses can be interchanged between combinations, and some poses may be mutated in. This process is repeated over a series of generations, converging on a final population that resembles the Pareto front of the brute force method. The Pareto front is found for this final population and the desired combination is selected using the Nash bargaining solution.

F. Validation with toy model

A simple toy model is created to validate the brute force approach and benchmark the performance of the GA. The toy model consists of a 10 m x 10 m x 10 m voxel grid with a 0.5 m voxel pitch. The

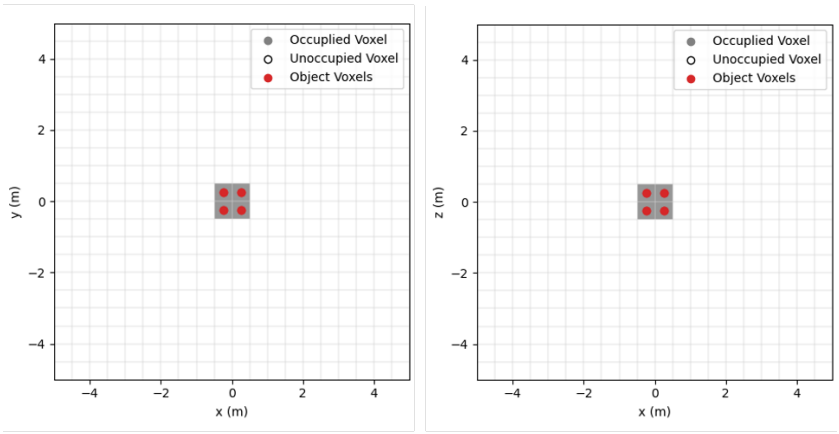

Fig. 2 Left: x-y plane of toy model Right: x-z plane of toy model

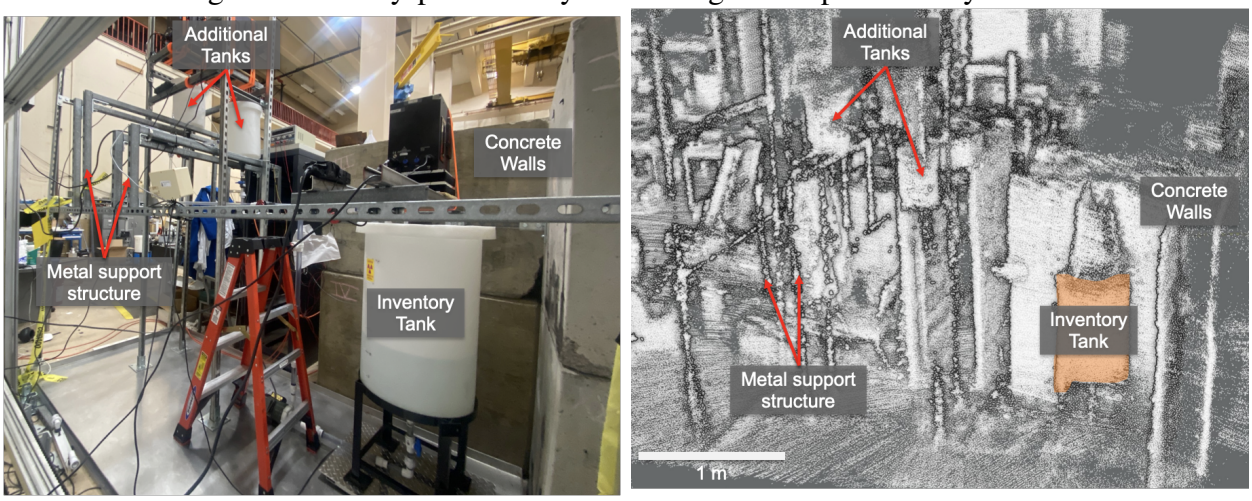

Fig. 3 Left: camera image of mock holdup loop. Right: 3D point cloud map of mock holdup loop, with the inventory tank highlighted in orange.

object of interest consists of 8 occupied source voxels in the center of the voxel grid. All other voxels in the scene are assumed to be unoccupied. Fig 2 shows the extent of the voxel grid with the volume of interest in the center.

G. Validation with safeguards-relevant measurement

After validation with the toy model, the GA-accelerated optimization is performed on a mock material holdup loop. The loop consists of two tanks, with an inventory tank where 54.4 μCi of F-18 is diluted in 32.2 L of water. Fig. 3 left shows a camera image of the mock holdup loop. The source is mixed in the water with a mixing loop attached to the inventory tank. The loop allows the water to be pumped between the inventory tank and the additional tanks, with the ability to divert material between the additional tanks to simulate material holdup. The H420-LAMP is used to generate a 3D map of the loop and surrounding environment, shown in Fig. 3 right. The inventory tank was selected within this 3D map as the volume of interest, and the optimal static measurement poses were determined with the GA approach. A 10-minute measurement is taken at each of the optimal poses, and quantitative Compton imaging is performed on the resulting data.

III. Results

A. Validation with toy model

The potential measurement positions are generated for the toy model with a constraint that all of the positions are on-axis with the object of interest in the z dimension with a maximum standoff distance of

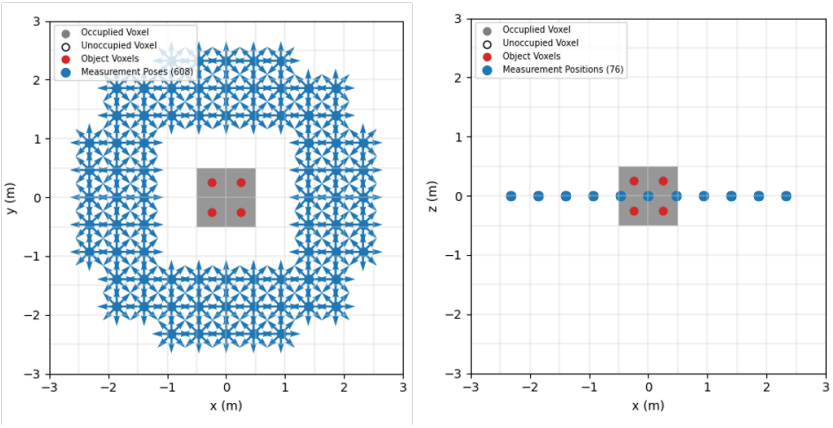

Fig. 4 Potential measurement poses selected for toy model shown in x-y plane (left) and x-z plane (right).

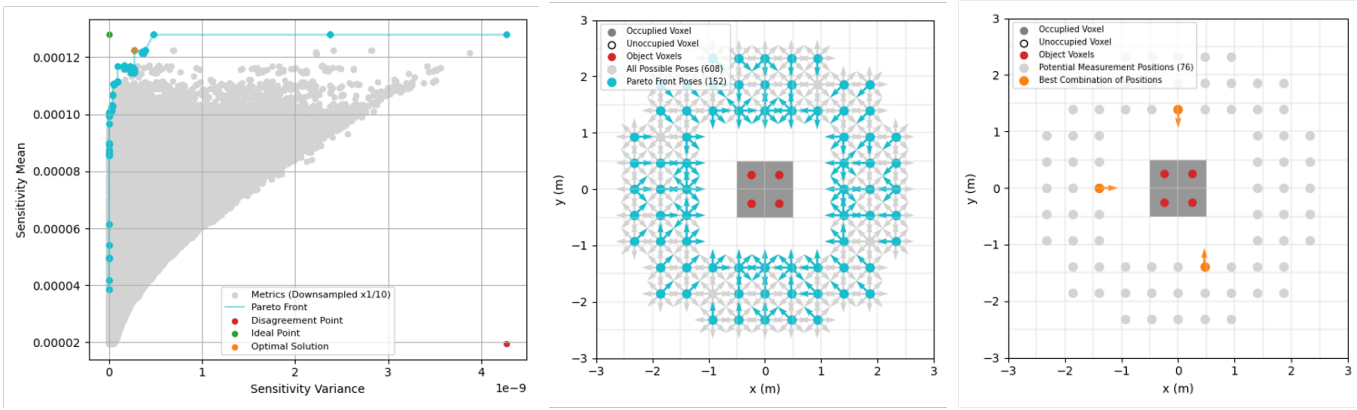

Fig. 5 Left: Optimization metrics of the combinations of positions and their Pareto frontier. Middle: Poses within combinations included in Pareto frontier shown in x-y plane. Right: Combination of positions identified by the Nash Bargaining Solution. Degenerate solutions exist in 90 ° rotations around the object of interest and in their respective mirrored configurations.

2.5 m from the center of the object, resulting in 76 x-y-z positions. The orientation is discretized in 45° increments around the z-axis, resulting in 8 potential orientations. Because the only occupied voxels in the toy model are the source voxels, the ray-tracing techniques do not mask out any of the measurement positions as infeasible. The 608 potential measurement poses generated are shown in the x-y and x-z plane in Fig. 4.

A combination of 3 measurement poses from the 608 candidate poses results in 43,895,680 potential combinations to perform the optimization over. The sensitivity to each of the 8 object voxels is computed for all 608 poses, and the metrics of interest are computed for each potential combination for the brute-force optimization. Fig. 5 left shows the sensitivity mean and variance metrics for a down-sampled fraction of all possible combinations of poses in grey. The cyan curve shows the Pareto front for these metrics, representing the set of optimal combinations for the toy model. Fig. 5 middle shows where the poses that comprise these optimal combinations lie in the x-y plane. The solution identified by the Nash bargaining solution is shown in orange in Fig. 5 left, lying close to the 'ideal point' of maximum sensitivity and minimum variance. The poses that comprise this combination are shown in Fig. 5 right. Due to the symmetrical nature of this toy model, there are degenerate combinations with the same metrics in 90^o rotations around the object and in mirrored configurations.

The GA is also used to find the optimal positions, and benchmarked against the brute-force method to demonstrate the feasibility of using this approach. Fig. 6, left shows the initial population of 1000 combinations used in the optimization, where 602/608 potential poses are represented. The initial pop-

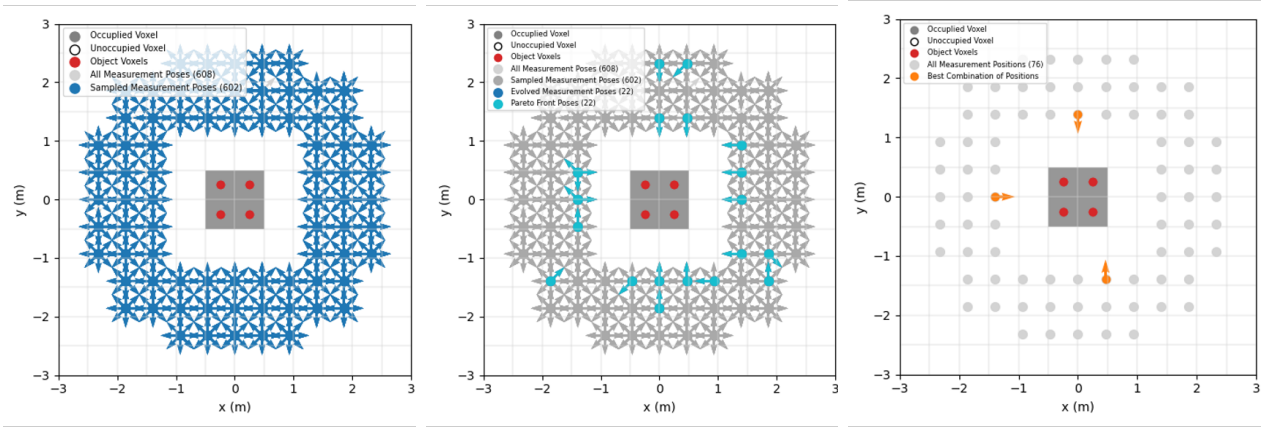

Fig. 6 Left: Initial population randomly sampled by the GA. Middle: Poses present in population after 50 evolutions and the Pareto frontier of the final population. Right: Combination of positions identified by the Nash Bargaining Solution. Degenerate solutions exist in 90 ° rotations around the object of interest and in their respective mirrored configurations.

ulation is evolved for 50 generations, with the default crossover parameter and 0.2 mutation parameter to allow for mutation of the 6 unrepresented positions into the population. Fig. 6 middle shows the final population on the x-y plane, where all of the poses lie on the Pareto front of the brute-force method as shown in Fig. 5 middle. Fig. 6 right shows the optimal combination of positions selected by the Nash bargaining solution, which matches the optimal combination of positions selected by the brute-force method (including the same degenerate solutions). The GA approach only required 50,000 computations (population of 1000 evaluated over 50 generations), while the brute-force approach required 43,895,680 computations, demonstrating that the same result can be reached with significantly fewer computations when using the genetic algorithm. This validates that this problem is a suitable use of the GA, facilitating the use of this optimization framework in the field.

B. Validation with safeguards-relevant measurement

The optimization approach is used to identify static measurement positions to measure the activity in the inventory tank shown in Fig. 3 right. Fig. 7 shows all of the potential measurement positions identified as feasible in grey. The measurement positions were constrained in z by the limits of the tripod that the imager rests on, and by a maximum standoff distance of 3 m. Each of these positions has line of sight to at least one of the inventory tank voxels as determined by the ray-tracing algorithm, and none of them reside in an occupied voxel. There are 8 potential orientations in 45° increments around the z-axis. The complex nature of this scene results in a fairly restricted number of potential measurement poses due to the voxel grid occupancy and ray-tracing constraints.

The GA optimization approach was used to find the optimal combination of poses. The GA used an initial population of 1000 combinations with the default crossover parameter and 0.2 mutation parameter. Fig. 7 shows the optimal poses identified by the Nash bargaining solution in orange, with a degenerate solution shown in blue. Data were collected at each of these poses for 10 minutes.

The Compton image of the inventory tank and its estimated activity are shown in Fig. 8. The source is localized to the position of the water within the inventory tank. The sum of activity in the source voxels is 7.8 μCi of F-18. A preliminary attenuation correction factor of 4.2 is applied to this reconstruction to yield a total source activity of 32.5 μCi.

IV. Discussion

The methods developed in this work present improved use of scene data to determine optimal measurement poses for quantitative gamma-ray imaging of volumetric sources. Identifying optimal measurement

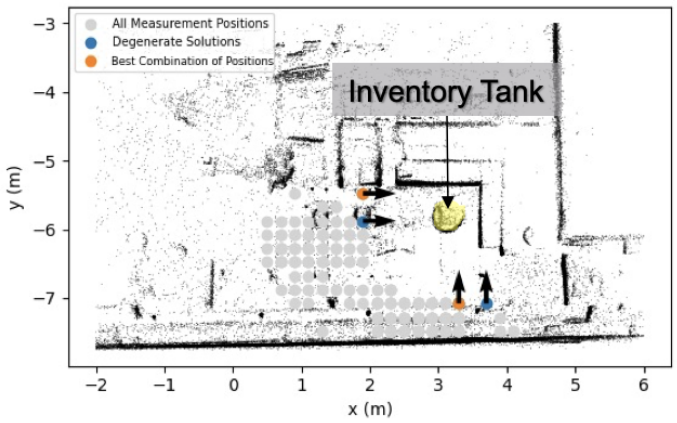

Fig. 7 Potential measurement positions identified for the inventory tank (highlighted in yellow), with the optimal combination of positions identified by the Nash bargaining solution and its degenerate solution identified.

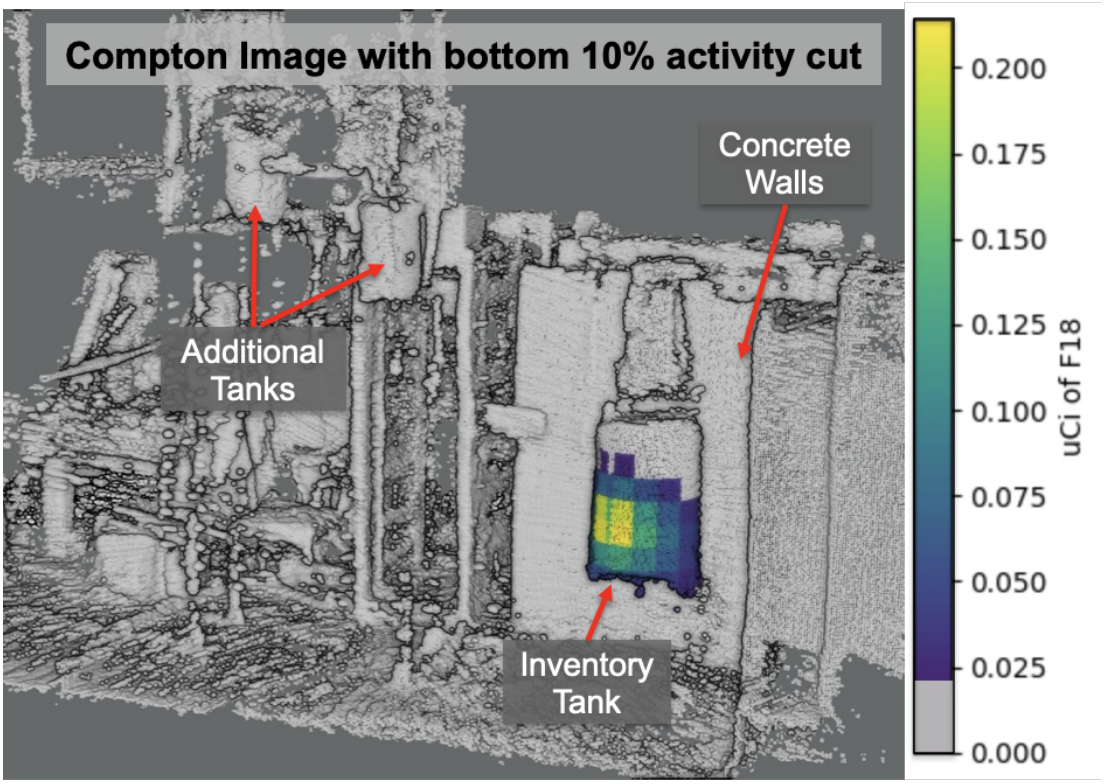

Fig. 8 Quantitative Compton image of inventory tank with bottom 10% of activity suppressed.

poses is useful in a variety of applications, including safeguards, where an analyst has limited time to make decisions about how to measure relevant materials. Placing the imager at poses with optimal sensitivity to the volume of interest permits shorter measurement times while still maintaining desired counting statistics, which results in reduced disruption to facility operation. Additionally, the point-and-dwell measurement mode employed in this work allows for reduced dose to personnel while measuring radiation data.

One drawback of using a GA is that there is not a guarantee to arrive at the same solution as the brute-force method. However, due to the nature of multi-objective optimization having multiple optimal solutions, it is very likely that the solution determined by a GA will still be a Pareto optimal solution. While the solution might not always be identical to the brute-force method, the time savings of using algorithmic approaches enables the use of these methods in real-time.

Another drawback of using the GA to solve this problem is that the parameters of the GA might

need to be altered in different scenarios, which would require operator knowledge of how the algorithm works. In the toy model example, 602 out of 608 measurement poses were represented in the initial sample. If significantly fewer measurement poses were represented, either the population size or the mutation parameter should be increased to more effectively represent the global population. However, if the population size is increased too much the computation time might become infeasible, and if the mutation parameter is set too high, the population may not converge. Algorithmic approaches other than the GA, such as gradient-free algorithms, could be investigated as additional means to solve this problem, which might require less parameter tweaking.

The LiDAR point cloud generated for the mock holdup loop contains pipes with a small diameter, which are removed by the filtering processes in the SLAM algorithm. Some fraction of source-containing water is in these pipes during the measurement, but because the pipes are not in the point cloud the image is not reconstructed in the pipes. If the filtering processes are removed from the SLAM algorithm, artifacts appear in the point cloud from movement of the operator in the scene, which is also undesirable. The shortcomings of LiDAR point clouds could be addressed with the additional sensors on the imaging system. Photogrammetry can be performed with the depth images returned by the depth camera, which could generate a higher fidelity point cloud.

The imaging algorithms used in this work do not yet include the ray-tracing estimates used for modeling attenuation in the optimization approach. Currently an attenuation factor computed from the source geometry is included after the image reconstruction to account for the source self-attenuation. The estimated activity in source voxels is expected to be closer to ground truth when the attenuation coefficient estimates are fully incorporated into the reconstruction algorithms, rather than applying a correction factor after the fact.

This work could also be improved by further leveraging the scene information to perform semantic segmentation techniques, which can identify the object of interest within a scene and inform attenuation estimates. Additionally, techniques could be developed to adaptively determine upcoming measurement poses as information is collected during measurements, rather than computing optimal poses from an initial scene survey alone. This approach could be useful in facilities where generation of a 3D map of the scene is prohibited or a very limited amount of time is available to take measurements.

V. Conclusions

In this work, we have shown preliminary results that demonstrate the feasibility of using scene data to determine optimal measurement poses for radiological measurements. A brute-force method is developed and tested on a toy model to consider all potential measurement poses in a scene, resulting in sensible measurement poses but requiring many computations. A genetic algorithm is used to accelerate this optimization, and benchmarked against the brute-force method, showing that it is a feasible acceleration to the problem that identifies the same optimal poses. These techniques are then applied to a mock holdup loop with a tank containing F-18 diluted in water. The optimal poses identified by the algorithm are used to measure the tank and generate a quantitative 3D Compton image from a 20 minute measurement. This technique presents a promising method to inform operators on optimal device operation in safeguards field measurements.

References

[1] K. Vetter *et al.*, "Advances in Nuclear Radiation Sensing: Enabling 3-D Gamma-Ray Vision," *Sensors*, vol. 19, no. 11, p. 2541, Jun 2019.
[2] C. Wahl *et al.*, "Coded-aperture imaging with high-resolution large-volume czt," in *2018 IEEE NSS/MIC*, 2018, pp. 1–5.
[3] R. Pavlovsky *et al.*, "3-D Radiation Mapping in Real-Time with the Localization and Mapping Platform LAMP from Unmanned Aerial Systems and Man-Portable Configurations," *arXiv:1901.05038*, 2018.
[4] W. Hess *et al.*, "Real-Time Loop Closure in 2D LIDAR SLAM," *in Proc. IEEE Int. Conf. Robot. Auto.*, pp. 1271–1278, 2016.
[5] R. Barnowski *et al.*, "Scene Data Fusion - Real-time Standoff Volumetric Gamma-ray Imaging," *Nucl. Instrum. Methods A*, vol. 800, pp. 65–69, 2015.
[6] A. Haefner *et al.*, "Handheld Real-time Volumetric 3-D Gamma-ray Imaging," *Nucl. Instrum. Methods A*, vol. 857, pp. 42–49, 2017.
[7] J. Hecla, K. Knecht *et al.*, "Polaris-lamp: Multi-modal 3-d image reconstruction with a commercial gamma-ray imager," *IEEE TNS*, vol. 68, no. 10, pp. 2539–2549, 2021.
[8] D. Hellfeld *et al.*, "Free-moving Quantitative Gamma-ray Imaging," *arXiv:2107.04080 [physics]*, Jul. 2021.

[9] M. S. Bandstra *et al.*, "Improved Gamma-Ray Point Source Quantification in Three Dimensions by Modeling Attenuation in the Scene," *IEEE TNS*, vol. 68, no. 11, pp. 2637–2646, Nov. 2021.
[10] J. Amanatides and A. Woo, "A fast voxel traversal algorithm for ray tracing," *Proceedings of EuroGraphics*, vol. 87, 08 1987.
[11] P. Vanderschraaf, "Learning bargaining conventions," *Social Philosophy and Policy*, vol. 35, no. 1, p. 237–263, 2018.
[12] J. Nash, "The bargaining problem," *The Econometric Society*, vol. 18, pp. 155–162, 04 1950.
[13] E. Kalai and M. Smorodinsky, "Other solutions to nash's bargaining problem," *The Econometric Society*, vol. 43, pp. 513–518, 05 1975.
[14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.
[15] F. Biscani and D. Izzo, "A parallel global multiobjective framework for optimization: pagmo," *Journal of Open Source Software*, vol. 5, no. 53, p. 2338, 2020. [Online]. Available: https://doi.org/10.21105/joss.02338