

Lawrence Berkeley National Laboratory

LBL Publications

Title

Surrogate distributed radiological sources I: point-source array design methods

Permalink

<https://escholarship.org/uc/item/8nr3401k>

Authors

Vavrek, Jayson R

Bandstra, Mark S

Hellfeld, Daniel

et al.

Publication Date

2024

DOI

10.1109/tns.2024.3351597

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Surrogate distributed radiological sources I: point-source array design methods

Jayson R. Vavrek, Mark S. Bandstra, Daniel Hellfeld, Brian J. Quiter, and Tenzing H.Y. Joshi

Abstract—In this first part of a multi-paper series, we demonstrate a method for using arrays of point sources to emulate continuous distributed gamma-ray sources when measured from a standoff of at least several meters. The method relies on the Poisson deviance statistic to test whether the array source “looks like” its continuous analogue when measured by a particular gamma-ray detector moving through 3D space on a particular trajectory. This point-source method offers significant advantages over truly distributed sources such as powders, solutions, or aerosols; notably, arrays of sealed point sources are safer to both personnel and the environment, and are more easily deployed, reconfigured, ground-truthed, and removed. We use this Poisson deviance metric to design eight different mock distributed sources, ranging in complexity from a 36×36 m uniform square grid of 5 mCi Cu-64 sources to a configuration where regions of higher and zero activity are superimposed on a uniform baseline. We then present several example calculations for various detector systems, altitudes, array source spacings, and source patterns, and examine under what parameters it is possible to design a surrogate array source that is nearly indistinguishable from a truly continuous distributed source. In Part II, we will detail the design, manufacture, and testing of Cu-64 sealed sources at the Washington State University research reactor, discuss their deployment during the aerial measurement campaign, and present results from several measurements.

Index Terms—gamma-ray imaging, distributed sources, airborne survey, Poisson deviance

I. INTRODUCTION

QUANTITATIVELY mapping continuous distributed radiological sources is important for radiological emergency response, whether the cause of the radioactive release is accidental (e.g., contamination from a reactor accident), or intentional (e.g., nuclear warfare). On short time scales, quantitative radiation maps can for instance inform which contaminated areas should be avoided by first responders while completing search and rescue or other missions, while on longer time scales, such maps can inform radiological remediation efforts. Testing and validating mapping and imaging algorithms for such distributed sources is challenging, however, for three related reasons. First, it is difficult to manufacture and deploy truly distributed radiation sources—radioactive material would have to be powdered, aerosolized, or dissolved, which can present a substantial human and environmental safety hazard [1], especially if the radioactive material were ingested or inhaled. For instance, measurements of dispersed activated KBr [2] or La_2O_3 [3] powder in ~ 1 Ci (37 GBq) quantities require substantial personal protective equipment (PPE) [4, Annex 3] and large standoff distances due

to high concentrations of airborne and deposited radioactivity. Second, powdered, aerosolized, or dissolved radionuclides cannot be easily and safely reconfigured, making it difficult to rapidly test multiple source configurations, to transfer the source material back to a laboratory for later assay, or—if the source half-life is long—to ensure that all radioactive material is removed from the environment. Third, producing truly continuous distributions of radioactive material with known ground truth patterns is also difficult—the deposition of powdered, aerosolized, or dissolved materials may deviate from the intended pattern due to factors such as changing winds or uneven mixing due to the mechanical variability of the depositor in inclement weather [1]. Post-deposition ground truth activity assays are possible using collimated high-purity germanium (HPGe) [5] or cerium bromide (CeBr_3) [6] detectors, but these measurements require close proximity to the source (increasing dose and often disturbing contaminated soil) and are limited to small ($\lesssim 1$ m²) areas in a single measurement, and thus are difficult to use for rapidly mapping large distributed sources spanning hundreds or thousands of square meters. While remotely- or autonomously-controlled ground robots could be used to carry the detectors used for the ground truth measurements and this would mitigate dose concerns, their use would create additional complications such as ensuring the robots did not disturb the source distributions or become radiologically contaminated themselves.

Instead, in this series of papers, we present a method for emulating truly continuous distributed radiological sources with arrays of sealed point sources, which are easily ground-truthable, re-configurable, and removable. The method is founded on the Poisson deviance statistic, and uses it in a hypothesis test to determine how much the array source “looks like” its continuous source analogue, for a given detector and trajectory. In presenting these calculations, we will rely on the concept of “spoofing”, wherein a successful spoof is one where the fake array of sources cannot be distinguished from the continuous source analogue. Using this metric as a guide, in this paper (henceforth Part I) we simulate unmanned aerial system (UAS) borne gamma-ray measurements (using the NGLAMP [7] and MiniPRISM [8] detectors) of eight different planar array source patterns of ~ 500 mCi of Cu-64 each. Then in Part II, we will detail the design and deployment of 300 ~ 7 mCi (~ 259 MBq) Cu-64 sources in these eight array source patterns during an August 2021 outdoor distributed sources measurement campaign at Washington State University (WSU). Additionally we will discuss the measured gamma-ray signatures and efforts to compare them against ground truth expectations. These measurements will in turn

JRV, MSB, DH, BJQ, and THYJ are with the Applied Nuclear Physics program at Lawrence Berkeley National Laboratory.

form the basis of an upcoming study comparing quantitative MAP-EM [9], [10] reconstructions of the source distributions against known ground truth.

II. METHODS

A. Mathematical framework

The degree to which a detector can distinguish a continuous source from an array depends on a number of factors, including the detector trajectory, the detector response, and the source itself. We consider a detector trajectory $\vec{\mathbf{r}} \in \mathbb{R}^{I \times 3}$ with dwell times $\mathbf{t} \in \mathbb{R}_+^I$, where I is the number of individual time-binned measurements made over the course of the full measurement. We note that we will use boldface italic to denote arbitrary vectors $\mathbf{v} \in \mathbb{R}^N$, vector arrow notation for a single 3D vector $\vec{v} \in \mathbb{R}^3$, and both to denote a collection of N vectors in 3D, $\vec{\mathbf{v}} \in \mathbb{R}^{N \times 3}$. For simplicity we also consider only a single detector, though this assumption can be readily generalized to J detector elements. In addition to its position vector, the detector also has a vector of orientations \mathbf{a} , each element of which can be described by a quaternion or rotation matrix. Each pair of position and orientation, known as a “pose”, in turn influences the detector response or “effective area” $\eta(\mathbf{a}, \vec{\mathbf{r}} - \vec{\mathbf{r}}')$ to a point source of radiation located at $\vec{\mathbf{r}}'$. The effective area is defined as the rate of particles detected per unit flux density at the detector, and is therefore proportional to both the intrinsic and geometric efficiencies: $\eta = 4\pi|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^2 \epsilon_{\text{int}} \epsilon_{\text{geo}}$ [11, footnote 1].

A truly continuous distributed source of radiation has a per-volume intensity distribution $w(\vec{\mathbf{r}}')$, $\forall \vec{\mathbf{r}}' \in \mathbb{R}^3$. The expected number of photopeak counts detected at each pose $\boldsymbol{\lambda} \in \mathbb{R}_+^I$ involves integrating over the source distribution,

$$\boldsymbol{\lambda} = \int_{\mathbb{R}^3} \frac{w(\vec{\mathbf{r}}') \boldsymbol{\eta}(\mathbf{a}, \vec{\mathbf{r}} - \vec{\mathbf{r}}') \mathbf{t}}{4\pi|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^2} e^{-\mu_{\text{air}}|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|} d^3 \vec{\mathbf{r}}', \quad (1)$$

where we have assumed that the only attenuation between the source and detector is from air, with an attenuation coefficient μ_{air} . Given that this work focuses on planar sources, we note that for an isotropic detector at a constant height $z = h$ above the center of a uniform circular plane source of radius R at $z = 0$ with an activity density $w(\vec{\mathbf{r}}') = \delta(z)w_0$, $0 \leq |\vec{\mathbf{r}}'| \leq R$ and no air attenuation, we have the analytical solution

$$\lambda_i = \frac{w_0 \eta_i t_i}{4} \log(1 + R^2/h^2), \quad i = 1, 2, \dots, I. \quad (2)$$

We note that Eq. 2 decreases more slowly with height h compared to the familiar $1/h^2$ behavior that occurs for point sources as well as when $h \gg R$.

For arbitrary distributions $w(\vec{\mathbf{r}}')$, however, the integral in Eq. 1 may be difficult or impossible to evaluate analytically (see, e.g., Refs. [12], [13, Appendix A], [14, Chapter 4]). A more computationally-oriented approach suitable for arbitrary distributions involves discretizing the distribution $w(\vec{\mathbf{r}}')$ into K point sources, in which case the expected number of counts is

$$\boldsymbol{\lambda} = \sum_{k=1}^K \frac{w_k \boldsymbol{\eta}(\mathbf{a}, \vec{\mathbf{r}} - \vec{\mathbf{r}}'_k) \mathbf{t}}{4\pi|\vec{\mathbf{r}} - \vec{\mathbf{r}}'_k|^2} e^{-\mu_{\text{air}}|\vec{\mathbf{r}} - \vec{\mathbf{r}}'_k|}. \quad (3)$$

Here the w_k would typically be chosen by voxelizing (a bounded subset of) \mathbb{R}^3 , evaluating $w(\vec{\mathbf{r}}')$ at each of the K voxel centers, and multiplying by each voxel volume. As K increases and the voxel size decreases (for a fixed total volume), the fidelity to Eq. 1 increases, but so too do the computational and storage costs. In either case, using Eq. 1 or Eq. 3, a specific realization of detected counts can then be generated by Poisson sampling the mean count vector $\boldsymbol{\lambda}$:

$$\mathbf{n} \sim \text{Poisson}(\boldsymbol{\lambda}). \quad (4)$$

Since Eq. 3 can also be used if the source is truly a collection of K individual point sources \mathbf{w} , it is our primary tool for computing expected count rates for both the array source \mathbf{w} of K points and a computational approximation to $w(\vec{\mathbf{r}}')$ of $K' \gg K$ points. Throughout this paper, we will therefore refer to both “continuous” and “array” sources as collections of discrete source points—where it is understood that the number of points K' in the former will be much larger than the number of points K in the latter—and when necessary will use the prefix “truly” if talking about the source distributions $w(\vec{\mathbf{r}}')$ in Eq. 1.

It is also useful to introduce the sensitivity map $\boldsymbol{\varsigma} \in \mathbb{R}_+^{K \times I}$, which is calculated for each potential source point as

$$\varsigma_k = \sum_{i=1}^I \frac{\eta(\mathbf{a}_i, \vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k) t_i}{4\pi|\vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k|^2} e^{-\mu_{\text{air}}|\vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k|}. \quad (5)$$

The sensitivity has dimensions of time, and thus can be interpreted as the expected number of counts per unit emission rate at the source point $\vec{\mathbf{r}}'_k$.

Finally, it is often useful to cast Eqs. 3 and 5 in matrix form by extracting the terms

$$V_{ik} \equiv \frac{\eta(\mathbf{a}_i, \vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k) t_i}{4\pi|\vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k|^2} e^{-\mu_{\text{air}}|\vec{\mathbf{r}}_i - \vec{\mathbf{r}}'_k|}. \quad (6)$$

into the “system matrix” $\mathbf{V} \in \mathbb{R}_+^{I \times K}$. Then the mean counts array is $\boldsymbol{\lambda} = \mathbf{V}\mathbf{w}$ and the sensitivity map is $\boldsymbol{\varsigma} = \mathbf{V}^T \mathbf{1}_I$, i.e., the sum of each column of \mathbf{V} .

B. Continuous source emulation

Since our primary consideration in this study is designing array sources that are relatively easy to deploy in the field, we consider only 2D regular grids of potential source locations (for both the array and continuous sources), with potentially multiple sources (i.e., variable activity) per location. We note however that non-planar fields and/or various support structures and fasteners could extend this method to 3D, and that irregular grids such as general tetragonal or cylindrical grids could be used (though may be more difficult to accurately deploy in the field). In turn, rather than starting with an arbitrary continuous source and seeking to create a representative array source, we consider the reverse problem: given an easy-to-deploy array source, what are the possible continuous sources from which it could have originated? As we are essentially up-sampling the spatial resolution of the source, there is no unique solution to this problem.

A solution that minimizes the amount of added information is to perform a nearest-neighbor interpolation from the array to

the continuous source points—i.e., initialize each interpolated continuous point activity to the activity at the nearest array point activity—and then uniformly scale the activities at the interpolated points so that their activity sum matches that of the array source. We ensure that each array point has an equal number of such interpolated points that are closest to it to minimize boundary effects. As a result, we must also extrapolate this “interpolation” beyond the array source boundary, but only to half as many points. Boundary effects could likely be further reduced by changing the density of points or by using a linear activity interpolation (instead of nearest-neighbor) near activity boundaries, but here we restrict ourselves to the constant nearest-neighbor method for simplicity.

The spacing of the interpolated continuous points is chosen based on two criteria. First, two neighboring array points must have an even number of interpolated continuous points between them, so that there is no middle point that would require activity fractionation when located at a non-uniform position within the array. Second, the continuous spacing must be much smaller than the array spacing. The interpolated continuous spacing required will depend on the array source spacing, as well as the detector altitude, trajectory, intrinsic efficiency, and angular resolution. As discussed in Section III, we designed raster pattern trajectories to cover both source and background areas of the field, with raster speeds and spacings determined by nominal UAS battery lifetimes. Flight altitudes were limited to between 5 and 15 m above ground level (AGL) due to ease of operation above 5 m and airspace restrictions above 15 m. The lowest flight altitudes then provide constraints on both the continuous and array source spacings. Using the methods described in Section II-D, we empirically find that an array spacing of 4 m and continuous spacings of $\lesssim 25$ cm are sufficient and computationally tractable forward models across our parameter space, given detector angular resolutions of $\sim 10^\circ$ for 511 keV singles for MiniPRISM and coarser for NG-LAMP. As shown in Fig. 1, using $m = 8$ interpolated continuous points on either side of an array point gives a continuous spacing of $4\text{ m}/(2 \cdot 8 + 1) = 0.235$ m, which is used throughout this work unless otherwise specified. The $m = 8$ continuous source is effectively a computationally feasible approximation to an $m \rightarrow \infty$ truly continuous distributed source, while the $m = 0$ array source is an experimentally feasible approximation to $m = 8$. The quality of these approximations can be determined using the statistical tools developed in Section II-D.

We note that in practice, setting the continuous point density too high can be computationally demanding and can induce roundoff error that appears to stem from summing large numbers of small contributions from weak source points with finite numerical precision. For the models in this paper, we find that above $m \sim 50$ continuous points per array half-spacing (below a ~ 4 cm continuous point spacing), the modeled count rates from the continuous source become unstable. Values of $m \simeq 8\text{--}50$ are stable and show minimal changes in expected counts, but these values will likely differ between scenarios. Using `float64` instead of `float32` precision may help eliminate this roundoff error, but our GPU-based

models (described in Section II-C) are currently limited to the latter.

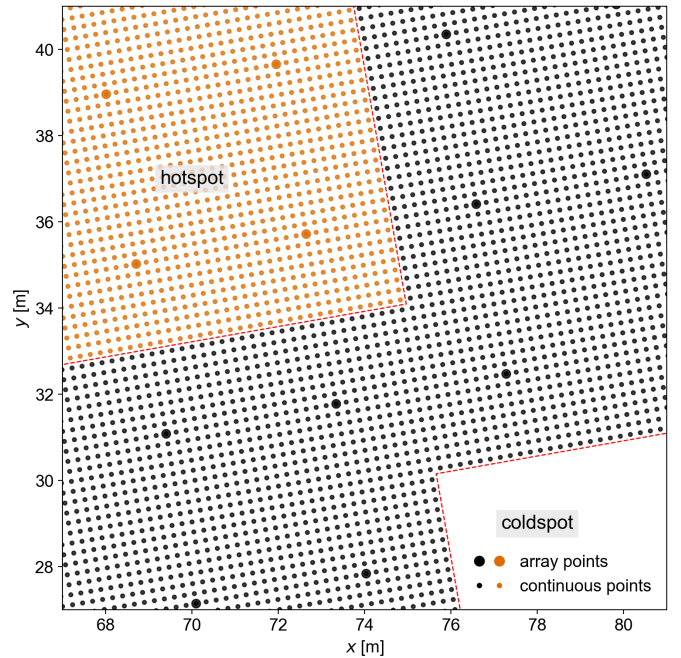


Fig. 1. Array (4 m spacing) and interpolated continuous ($4\text{ m}/17 = 0.235$ m spacing) sources for the hot/coldspot pattern discussed in Section III. The view is zoomed to more clearly show the design of the continuous source, especially at activity boundaries. The full array source is shown in Fig. 3.

C. Forward projections and detectors

The forward projections of both the continuous and array sources to each detector (accounting for their individual position offsets) are computed at each pose of the trajectory using Eq. 3. The forward projection is implemented in the Python-based, GPU-accelerated `mfd` (multi-modal free-moving data fusion) library [15], which computes an array of expected counts $\lambda \in \mathbb{R}_+^{I \times J}$, where I is the number of individual measurements and J is the number of individual detector elements. This work leverages the NG-LAMP [7] and MiniPRISM [8] detection systems, with $J = 4$ CLLBC crystals (of size $2 \times 2 \times 1$ inch) and $J = 58$ CZT crystals (of size $1 \times 1 \times 1$ cm), respectively.

Detector response functions were taken from existing characterizations of the NG-LAMP and MiniPRISM detectors, which were computed ahead of time using the Geant4 framework [16], [17], [18]. The detector systems (and Geant4 models thereof) include the Localization and Mapping Platform (LAMP), which consists of a LiDAR and an inertial measurement unit (IMU)—enabling LiDAR-based Simultaneous Localization and Mapping (SLAM) [19], [20], [21]—as well as a video camera, single-board computer, and front-end detector electronics—see Fig. 2. When coupled to a UAS, the systems can also read out real-time kinematic (RTK) and standard GPS positioning measured by the UAS as a comparison for LiDAR SLAM-computed positions.

Responses were computed for both single- and double-crystal (i.e., Compton) full-energy detection efficiency at




Fig. 2. Left: The NG-LAMP detector system, comprising four CLLBC crystals, a LiDAR unit, and a LAMP contextual sensor suite. Photo from Ref. [22]. Right: The MiniPRISM detector system, comprising 58 CZT crystals, a LiDAR unit, and a LAMP contextual sensor suite. Photo from Ref. [23].

511 keV, though the forward projections in this work model only the single-crystal signals. We additionally include a small but non-zero background rate in the photopeak region. In the simulation studies discussed in Section III we used 2 and 0.2 counts/s/detector for NG-LAMP and MiniPRISM, respectively, which were order-of-magnitude values estimated from previous indoor measurements. In the experimental analysis of Part II, however, we use 0.5 and 0.034 counts/s/detector, respectively, (i.e., 2 counts/s after summing over detectors) as determined by dedicated handheld and aerial background measurements at the WSU field.

D. Statistical tests

In evaluating the quality of an array source as a surrogate for a given continuous source distribution (and detector and detector trajectory), it is useful to distinguish whether one is approaching the problem as an anomaly *detection* or an anomaly *identification* problem.

In the case of anomaly detection, the question of whether the data are consistent with the given continuous source model naturally leads to the use of a goodness-of-fit statistic, which for Poisson-distributed data is the deviance or log-likelihood statistic [24]. For anomaly identification, the anomalous alternate model (here, the array source) is fully specified, and one would decide whether the array source or continuous source is a better fit to the data. This approach leads to the use of the likelihood ratio test (LRT) to determine which hypothesis (i.e., null/continuous or alternate/array) is more consistent with the data. The latter statistical test would in general be more powerful than the former because the analyst has complete information about the possible anomaly, and thus the required similarity between the continuous and array sources would be more stringent than in the anomaly detection case, which may be desired in some applications.

In this work, we adopt the anomaly detection approach. We are primarily interested in creating array sources that look reasonably like continuous sources, such that a data analyst would rarely reject the continuous source and could proceed along in using the data for various continuous source imaging applications without needing to know about the properties or even existence of the array source.

The aforementioned deviance $D(\mathbf{n}|\boldsymbol{\lambda})$ is a scalar goodness-of-fit metric for comparing a vector of observed counts \mathbf{n} to a

model $\boldsymbol{\lambda}$ of Poisson mean counts [24], [25], [26], and is given in “likelihood form” as

$$D(\mathbf{n}|\boldsymbol{\lambda}) = -2 \log L(\mathbf{n}|\boldsymbol{\lambda}) + 2 \log L(\mathbf{n}|\mathbf{n}), \quad (7)$$

where $L(\mathbf{n}|\boldsymbol{\lambda})$ is the Poisson likelihood of observing the data \mathbf{n} given the model $\boldsymbol{\lambda}$, and $L(\mathbf{n}|\mathbf{n})$ is the likelihood of observing the data if the model perfectly matched the data. Noting that $L(n_i|\lambda_i) = e^{-\lambda_i} \lambda_i^{n_i} / n_i!$ for a single measurement i and that likelihoods multiply since the measurements are statistically independent, we have

$$D(\mathbf{n}|\boldsymbol{\lambda}) = 2 \sum_{i=1}^I [n_i \log(n_i) - n_i \log(\lambda_i) + \lambda_i - n_i]. \quad (8)$$

Defining

$$p_i \equiv \frac{\lambda_i}{\sum_i \lambda_i} \equiv \frac{\lambda_i}{\bar{N}} \quad (9)$$

$$q_i \equiv \frac{n_i}{\sum_i n_i} \equiv \frac{n_i}{N} \quad (10)$$

we can re-write the deviance in the “magnitude and shape form” as

$$D(\mathbf{n}|\boldsymbol{\lambda}) = D(N|\bar{N}) + 2ND_{\text{KL}}(\mathbf{q} \parallel \mathbf{p}), \quad (11)$$

where $D(N|\bar{N})$ is the deviance between the observed sum of counts N and the model sum of counts \bar{N} , and $D_{\text{KL}}(\mathbf{q} \parallel \mathbf{p})$ is the Kullback–Leibler divergence between the normalized count vectors \mathbf{q} and \mathbf{p} . Note that in the case of multiple detector elements ($J > 1$), \mathbf{q} and \mathbf{p} are formed from concatenating the \mathbf{q} and \mathbf{p} of each separate detector. In the discussion below, we assume $J = 1$ for simplicity, but if $J > 1$ one can simply replace $I \leftarrow IJ$.

For multiple Poisson noise realizations of the I measurements, the distribution of the deviance statistic can be approximated by a shifted Gamma distribution whose first three moments match the calculated moments of the deviance statistic. The Gamma distribution is chosen to extend the moment-matching technique from a symmetric Gaussian distribution to one with a non-zero skew term. The parameters of this shifted Gamma distribution are expensive to calculate as the number of measurements I gets large, but two limiting cases exist when the number of counts n_i in most measurements is $\gtrsim 30$:

- 1) As the gross counts N become large, the deviance distribution can be well-approximated by a χ^2 distribution with I degrees of freedom.
- 2) As the number of measurements I becomes large, the χ^2 distribution itself can be well-approximated by a normal distribution with mean I and variance $2I$.

These limits will often not apply due to the low counts per measurement far from the source, so in general we will use the Gamma distributions. As shown in the later Fig. 7, however, the deviance distributions are often still well-described by normal distributions, but with means and variances different from I and $2I$.

We can now define the model $\boldsymbol{\lambda}_0$ as the mean count vector for our continuous source, and the model $\boldsymbol{\lambda}_1$ for our coarsely-gridded array source. Measurements \mathbf{n}_0 and \mathbf{n}_1 of the continuous and array sources will produce deviances of $D(\mathbf{n}_0|\boldsymbol{\lambda}_0)$

and $D(\mathbf{n}_1|\lambda_1)$, respectively. (Note that in the aforementioned large N and large I limits, both $D(\mathbf{n}_0|\lambda_0)$ and $D(\mathbf{n}_1|\lambda_1)$ will converge to the same distribution.)

We now know in principle what the deviance distributions will look like, given that we know which is the true model for a particular measurement, λ_0 or λ_1 . We denote their probability density functions (pdfs) as $\mathbb{P}(D|\lambda_0)$ and $\mathbb{P}(D|\lambda_1)$ for true models λ_0 and λ_1 , respectively. If we do not know the true model, however, we can ask what the deviance distribution will look like when the model is mis-specified. In particular, we can ask how the deviances $D(\mathbf{n}_1|\lambda_0)$ are distributed when the data are generated from an array source (\mathbf{n}_1) but deviances are calculated assuming a continuous source (λ_0).

We can then compute the theoretical shifted Gamma distributions of the deviances assuming the data samples \mathbf{n}_0 and \mathbf{n}_1 are Poisson samples from λ_0 and λ_1 . In particular we compute four theoretical deviance distributions for each parameter combination (see for instance the later Fig. 6):

- 1) $D(\mathbf{n}_0|\lambda_0)$: the deviance distribution when Poisson samples \mathbf{n}_0 are generated from the continuous source model λ_0
- 2) $D(\mathbf{n}_1|\lambda_1)$: the deviance distribution when Poisson samples \mathbf{n}_1 are generated from the array source model λ_1
- 3) $D(\mathbf{n}_1|\lambda_0)$: the deviance distribution when Poisson samples \mathbf{n}_1 are generated from the array source model λ_1 but their deviances are calculated assuming the continuous model λ_0 is the correct model.
- 4) $D(\lambda_1|\lambda_0)$: a constant cross term that originates from assuming the incorrect model; this is not a true deviance as it does not compare Poisson counts to a model, but it does have the same functional form.

The shifted Gamma distribution of the “non-central” deviances $D(\mathbf{n}_1|\lambda_0)$ is computed from the first three moments in essentially the same fashion as the central deviances. Namely, since each deviance statistic is the sum of statistically independent terms, the mean, variance, and third central moment of the deviance are simply the sums of those same moments for the terms. Those three moments are numerically estimated using the relevant Poisson distribution. For example, following Eq. 8, the i^{th} term (the “unit deviance”) of $D(\mathbf{n}_1|\lambda_0)$ is

$$2(n_i \log n_i - n_i \log \lambda_{0i} + \lambda_{0i} - n_i) \quad (12)$$

and the mean, variance, and third central moment of this term can be estimated assuming $n_i \sim \text{Poisson}(\lambda_{1i})$. In the aforementioned large- N and large- I limits, these calculations also allow one to compute the theoretical parameters of the χ^2 or Gaussian approximations for $D(\mathbf{n}_1|\lambda_0)$.

Equipped with this framework, we can now describe how well the array source mimics the continuous source. Quantitatively, given the continuous and array models λ_0 and λ_1 , what is the probability that a sample \mathbf{n}_1 drawn from λ_1 “looks like” it was drawn from λ_0 ? I.e., what is the false negative probability P_{FN} for incorrectly deciding that the most likely model $\hat{\lambda}$ is the continuous source λ_0 when the true model is the array source λ_1 ? In the absence of prior information,

the decision rule for choosing the most likely model $\hat{\lambda}$ for an observed deviance D is simply

$$\hat{\lambda} = \begin{cases} \lambda_0, & D \leq D^* \\ \lambda_1, & D \geq D^* \end{cases} \quad (13)$$

where the decision threshold D^* is the value of D at which both models are equally likely:

$$D^* = \underset{D}{\text{argwhere}} \mathbb{P}(D|\lambda_0) = \mathbb{P}(D|\lambda_1). \quad (14)$$

Then the false negative probability is

$$P_{\text{FN}} = \mathbb{P}(\hat{\lambda} = \lambda_0 | \lambda = \lambda_1) \quad (15)$$

$$= \mathbb{P}(D < D^* | \lambda = \lambda_1) \quad (16)$$

$$= \text{CDF}_{\lambda_1}(D^*), \quad (17)$$

where CDF_{λ_1} is the cumulative distribution function of D given $\lambda = \lambda_1$. As expected, the probability that a sample \mathbf{n}_1 drawn from the array source λ_1 “looks like” it was drawn from a continuous source λ_0 —and thus the degree to which the array source can be used as a useful proxy of a continuous source—depends on the overlap of the two deviance pdfs. Intuitively, a perfect “spoof” should have $P_{\text{FN}} = 1/2$ —it is indistinguishable from the continuous source via the deviance metric. In this work, we relax this perfect spoof condition and consider an array source to be a *practical* spoof of a continuous source if it has $0.4 \leq P_{\text{FN}} \leq 0.5$, though we note that this lower bound is somewhat arbitrary.

III. DEMONSTRATION DESIGN

Eight planar array sources consisting of up to 100 individual 5 mCi Cu-64 sources each were designed for use during the measurement campaign of Part II. These configurations are, as shown in Fig. 3:

- 1) a 10×10 square;
- 2) a pair of 5×10 rectangles separated by two grid spacings (8 m);
- 3) a pair of 5×10 rectangles separated by three grid spacings (12 m);
- 4) a 3×16 line with a hot center, which has four sources per point, referred to as the “hot line”; and
- 5) an L-shape with 12 (13) sources on its short (long) dimension and a thickness of five sources;
- 6) a 9×9 square with inset 3×3 grids of four sources per point and no sources, referred to as the “hot/coldspot”;
- 7) a 5×9 rectangle with a gradient in intensity that is produced by three adjacent 5×3 rectangles containing three, two and one source per point;
- 8) a “plume” consisting of a 7×7 outer checkerboard pattern and a 5×5 fully-occupied central region.

As shown in Fig. 3, with the exception of the plume source, the source spacing in all source patterns is 4 m. The 4 m separation was chosen to provide good P_{FN} values while also creating spatially large source distributions.

The sources share a common lower-left corner at $(x, y) = (65, 10)$ m, chosen to reduce dose rates to personnel at the headquarters (HQ) and create large empty regions on the west side of the field (see Fig. 4) to enhance contrast between source

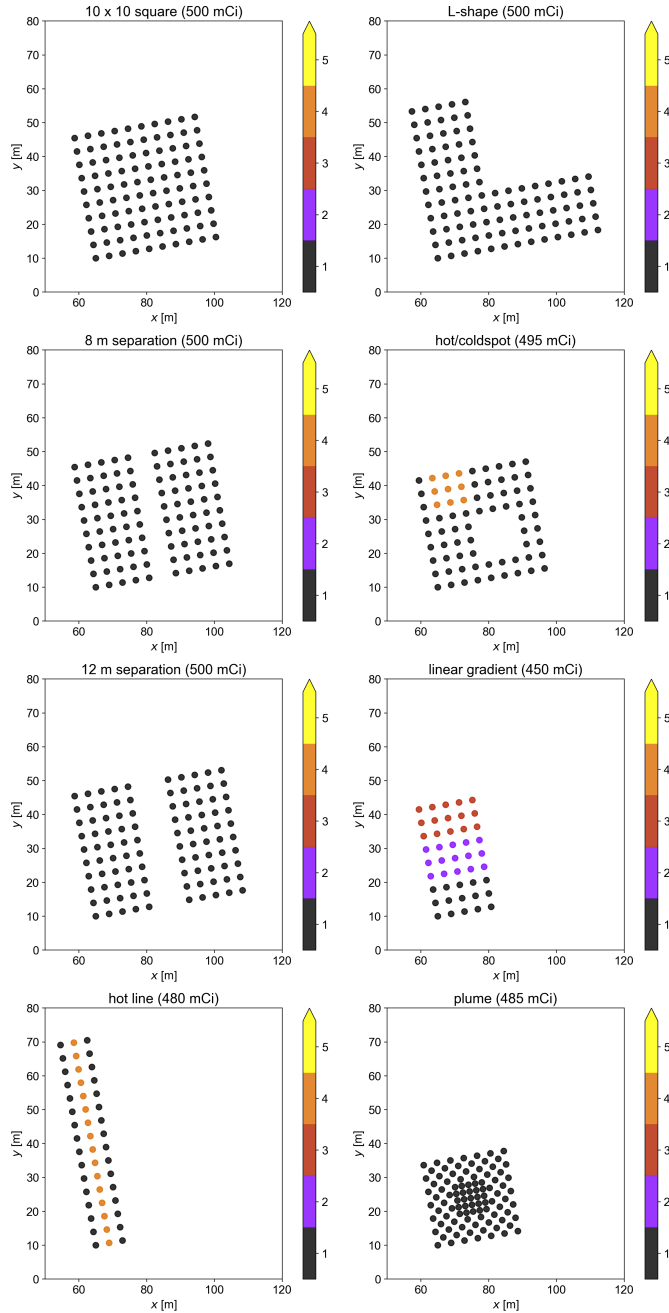


Fig. 3. All eight source configurations designed for the WSU measurement campaign. The color bar denotes the number of nominal 5 mCi Cu-64 sources at each point.

and background regions. Each array source pattern is rotated 10° about its lower-left corner to reduce the effect of aliasing with the UAS raster pattern, which is aligned with the field boundaries.

The above sources were designed for a variety of measurement goals, the analysis for which will be covered in a later work. The 10×10 square, being symmetric and uniform, provides a simple baseline with which to test reconstruction quality (via metrics such as the Structural Similarity Index Metric (SSIM), root-mean-squared error (RMSE), total activity, uniformity of activity, and edge sharpness) as well as

repeatability over multiple measurements. The L-shape was similarly designed as a uniform source with an interior corner. We note that a modified version of our L-shape could provide an array source analogue to the truly continuous source in Ref. [27]. The separated pairs of 5×10 rectangles were designed as uniform sources with narrow corridors of zero activity; such patterns can be analyzed for how well-separated the two rectangles are after reconstruction as a function of measurement parameters (e.g., altitude, detector, and the rectangle separation itself) and provides opportunities for studying dose-minimizing or information-maximizing path-planning algorithms. The linear gradient source was designed as a simple non-uniform source with which to test the reconstructed activity dropoff. Similarly, the hot/coldspot source provides non-uniform activities as well as internal regions of zero activity; this source will be used to again test the “sharpness” of the reconstruction. The hot line source was designed as a high-contrast source and also to test the resolution for narrow shapes with a spatial gradient. Finally, the plume source was also designed to test the activity change between the outer and inner regions, while also testing denser source placements.

The false negative probability P_{FN} is also influenced by the detector trajectory, and thus the trajectories flown by the UAS were also designed with a number of competing goals in mind. First, a dense field-aligned raster pattern was chosen to reduce the parameter space, avoid obstacles adjacent to the field, generate an approximately uniform sensitivity in the source region, and provide opportunities for studies of sparser raster patterns (e.g., only every second raster line) by cutting measurements in post-processing rather than by retaking data. As discussed in Section II-B, flight altitudes were limited to between 5 and 15 m. The UAS orientation was course-aligned in order to improve the LiDAR coverage vs. a fixed orientation. The raster spacing of 5.2 m, length of 100 m, and speed of 2.6 m/s were chosen to both overfly the source extent and collect data over zero-source regions for improved background estimation while completing in $\lesssim 10$ min to make the most use of the battery life. Originally, the raster was designed to traverse the entire long dimension of the field (up to the 5 m buffer on each side), but was shortened during the measurement campaign to approximately the 100 m width shown in Fig. 4 to account for lower-than-expected UAS battery performance. Similarly, most raster patterns were started from the bottom left corner of the field so that the passes over the source would be completed first in case of an early landing. Moreover, it was found that that this trajectory typically led to acceptable false negative probabilities of $0.4 \leq P_{FN} \leq 0.5$, and to generally accurate MAP-EM reconstructions of the simulated source shape and intensity from the forward-projected λ . We hypothesize this agreement between the simulated true and reconstructed sources is due in large part to the relatively high and uniform sensitivity (Eq. 5) over the true source extent afforded by this trajectory across various altitudes, including the 6 m altitude (AGL) used for most NG-LAMP flights—see again Fig. 4. In particular, using an altitude larger than the pass spacing tends to smooth out variations in the sensitivity due to the large constant z term in the distance between the detector and a given source voxel.

IV. RESULTS

Fig. 4 shows an example scenario consisting of a 10×10 square array source pattern (with 4 m source spacing) and a 6 m (AGL) NG-LAMP raster trajectory—see Section III for additional information on the design of the source and raster patterns. Fig. 5 shows the expected counts measured by the NG-LAMP detector with both the array (λ_1) and continuous (λ_0) sources. Here there are $I = 1043$ time bins of width $t_i = 0.5$ s and $J = 4$ detectors for a total of $IJ = 4172$ measurements. Only small deviations ($<0.2\sigma$ absolute) between the array and continuous source count vs. time profiles are observed. As visible here and in the later Fig. 10, slight edge effects are visible—the array source count prediction switches rapidly from a slight underestimate to a slight overestimate just outside vs just inside the source boundary due to differences in the positions of the outermost points in each source. In the deeper, more uniform parts of the sources, these discrepancies are reduced.

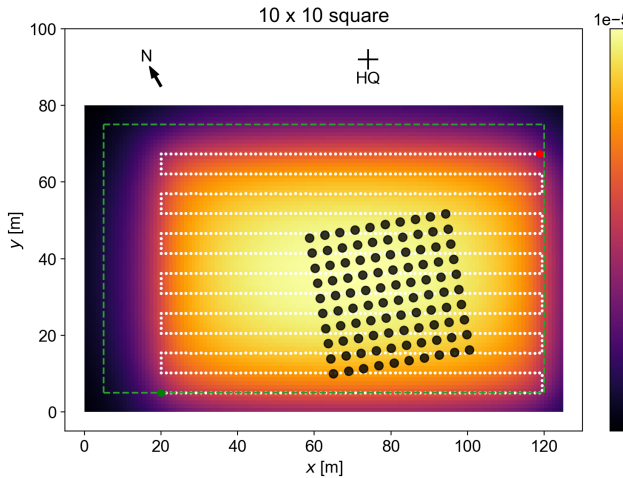


Fig. 4. Layout of the synthetic 10×10 square source (black points) and a typical 100 m-wide synthetic UAS raster pattern (dotted white line) with 13 lines spaced at 5.2 m. The start and stop of the trajectory are indicated by green and red circles, respectively. The green dashed rectangle denotes a 5 m flight buffer from the field boundary. The approximate compass direction and HQ location are also denoted. The sensitivity map ζ at 6 m (AGL) is shown beneath the source points and ranges from ~ 0.5 – 4.5×10^{-5} counts-s over the field extent. Over the source extent, the dense raster spacing maintains a relatively uniform sensitivity of ~ 3.5 – 4.5×10^{-5} counts-s.

Fig. 6 shows the theoretical and empirical distributions of deviance statistics computed from the two λ arrays (of length IJ , i.e., not summed over detectors) after 5000 Monte Carlo samples, and Fig. 7 shows the comparison of the deviance distributions $D(\mathbf{n}_0|\lambda_0)$ and $D(\mathbf{n}_1|\lambda_0)$ used for computing the false negative probability P_{FN} . For this combination of detector, trajectory, and source pattern, we find a false negative probability of $P_{\text{FN}} = 0.489$, almost indistinguishable from chance, indicating that the surrogate array source is of high fidelity.

The Poisson random samples of Fig. 6 also motivate the Monte Carlo study of other non-Poisson noise sources that may be present in real experiments but cannot easily be incorporated into the analytical models of Section II-D. In




Fig. 5. Top: forward projections of the synthetic mean counts of the continuous source (λ_0) and the array source (λ_1) for the scenario in Fig. 4. The “continuous” curve in blue is almost completely obscured by the “array” curve in orange. Bottom: comparison of the two λ , showing statistical differences (in this case, generally below 0.2σ absolute value) in the array source (and its Poisson error) compared to the continuous source. Counts are binned to $t_i = 0.5$ s and summed over the $J = 4$ NG-LAMP detector elements.

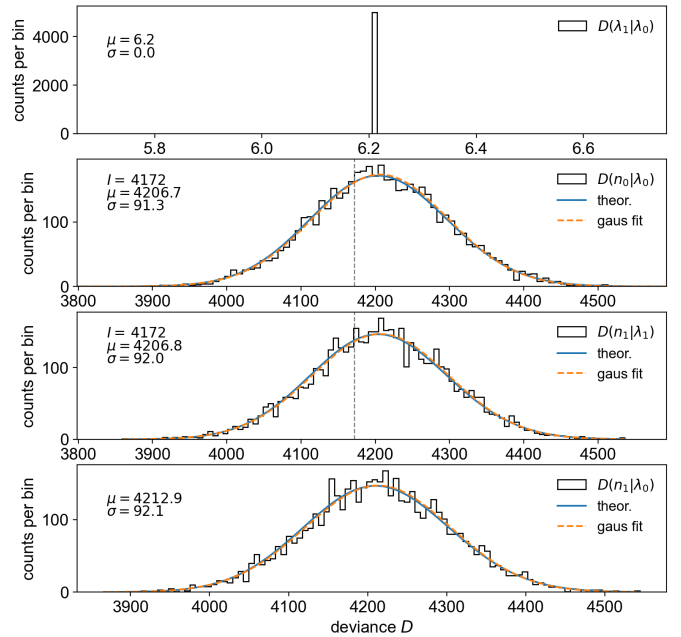


Fig. 6. Deviance distributions for the 6 m (AGL) NG-LAMP raster pattern over the 10×10 square of 5 mCi sources shown in Fig. 4. The histograms show the result of Poisson sampling λ 5000 times, while the solid and dashed curves show the theoretical shifted Gamma distributions and Gaussian fits to the histogram, respectively. In the top plot, the standard deviation σ is 0 since only the two mean count arrays are compared—since only Poisson noise is present, the $D(\lambda_1|\lambda_0)$ is constant. In the middle two plots, a dashed vertical line is drawn at $D = IJ$, the number of measurements, showing that distributions are approximately but not exactly centered at IJ . Top: $D(\lambda_1|\lambda_0)$ for the two models. Top middle: $D(\mathbf{n}_0|\lambda_0)$ for the continuous source. Bottom middle: $D(\mathbf{n}_1|\lambda_1)$ for the array source. Bottom: $D(\mathbf{n}_1|\lambda_0)$ for the array source compared to the continuous source.




Fig. 7. Deviances from Fig. 6 used to determine the false negative probability $P_{FN} = 0.489$. $\Delta\mu$ denotes the difference in means of the two distributions.

reality, a deployed array source will likely have some per-point activity and position noise that will further increase its deviance from the idealized continuous source. In Fig. 8, we inject a per-point Gaussian activity noise of $\pm 10\%$ (standard deviation) to the uniform square array source in the scenario of Fig. 6, and generate 5000 independent noised array source realizations. Because the λ_1 are now no longer constant, the $D(\lambda_1|\lambda_0)$ cross term is also no longer constant, and in fact contributes a substantial and non-Gaussian amount of deviance to the final $D(n_1|\lambda_0)$ distribution used to calculate P_{FN} . Fig. 9 shows how this excess or model-mismatch deviance grows with the magnitude of the activity noise. Conversely, we find that injecting a per-point Gaussian position noise of ± 0.05 m (standard deviation) to the array source introduces little extra deviance, merely broadening the $D(\lambda_1|\lambda_0)$ by about 0.5 (standard deviation). In realistic measurement scenarios, reducing source activity errors is therefore much more crucial than achieving precise spatial alignment in the array source.

We note that Fig. 4 also serves to introduce the Cartesian “field coordinate system” used throughout this work and Paper II. The origin is placed at the southwest corner of the field, and x and y are measured along the fence lines to the southeast and northwest corners, respectively. The positive y direction differs from north by roughly 30° . The z direction therefore defines elevation above the ground level, which is assumed to be perfectly flat.

Fig. 10 shows a second example scenario in which the same NG-LAMP detector system performs a 5 m (AGL) raster pattern over the hot/coldspot source shown in Fig. 3. Instead of showing differences between the array and continuous source counts in terms of σ vs time as in Fig. 5, in the top panel of Fig. 10 we plot the relative difference (summed over detectors and normalized to the continuous counts) as a function of xy position over the source. As expected, the array counts λ_1 are larger than the continuous counts λ_0 over the array source points and especially the hotspot, while the reverse is true just outside the source boundary and over the coldspot. The middle plot of Fig. 10 then shows the deviance accumulate over the measurement duration, with the largest increases occurring over the hotspot. The bottom plot shows the resulting deviance distributions, indicating a false negative probability of $P_{FN} = 0.452$. The separation in means between the two curves corresponds to the final accumulated deviance




Fig. 8. Deviance distributions for the scenario of Fig. 6 with an additional $\pm 10\%$ (standard deviation) per-point activity noise in the array source. Note that while the $D(n_1|\lambda_0)$ distribution (bottom) is no longer well-described by a Gaussian profile, its empirical mean and standard deviation are well-described by the linear and quadrature sums of the mean and standard deviation of the $D(\lambda_1|\lambda_0)$ (top) and $D(n_i|\lambda_i)$, $i = 0, 1$ (middle) distributions.




Fig. 9. Empirical means (points) and standard deviations (error bars) of the excess or model-mismatch deviance $D(\lambda_1|\lambda_0)$ as a function of the Gaussian activity noise magnitude for the 6 m (AGL) NG-LAMP raster over the 10×10 uniform square source. The point at 10% corresponds to Fig. 8.

in the middle plot, and the reduced P_{FN} compared to Fig. 7 is largely due to the 1 m reduction in altitude.

In Fig. 11, we show the variation of P_{FN} with array source spacings from 1–6 m at altitudes of 2–10 m (AGL) for the NG-LAMP detector system and the hot/coldspot source of Fig. 10 with baseline per-point activities of 5 and 1 mCi. The analogous continuous source spacings were also varied, holding constant the number of continuous points per array point. In this case, we arbitrarily used $m = 20$ continuous points in on each side of each array point instead of the $m = 8$ shown in Fig. 1, but the difference from $m = 8$ is minimal, once again indicating that $m = 8$ is an accurate and more computationally feasible approximation. We note that the bumps in the P_{FN} appear to be due to aliasing between

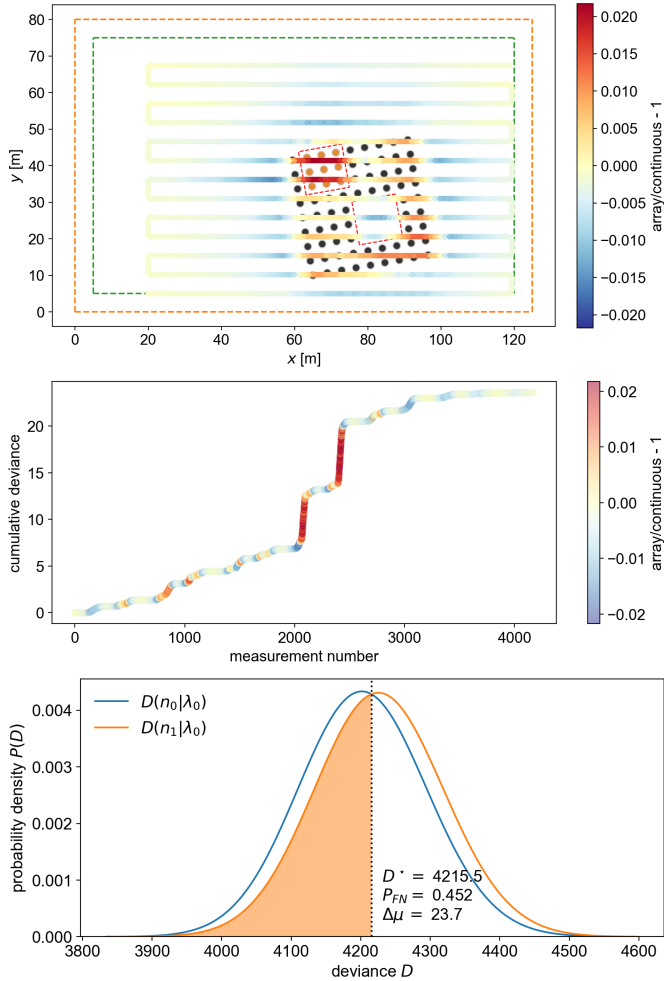


Fig. 10. Top: relative difference between the array and continuous source counts (summed over detectors and normalized to the continuous counts) as a function of detector position over the hot/coldspot source during a 5 m (AGL) NG-LAMP raster. The source points are colored by activity according to the color scale in Fig. 3. Middle: accumulation of deviance vs measurement number, with each point colored according to its relative difference as in the top plot. Bottom: resulting decision information plot, with a false negative probability of $P_{\text{FN}} = 0.452$. $\Delta\mu$ denotes the difference in means of the two distributions.

the trajectory pass spacing and the non-uniformity of the hot/coldspot source, as they largely disappear when using a uniform square source. Additionally, if the NG-LAMP detector response is replaced with an isotropic response function with the same mean effective area, the bumps in Fig. 11 change only slightly in magnitude and do not disappear. This is consistent with the expectation that the $1/r^2$ variation from the xy motion in such a wide-area raster pattern has a larger influence on the predicted counts λ than anisotropy in the response function. At 5 mCi per point, the array source approximation typically begins to fail ($P_{\text{FN}} \lesssim 0.40$) when the array source spacing approaches or exceeds the detector altitude. At 1 mCi per point, the P_{FN} increase due to reduced count rates. Figures such as Fig. 11 were used to choose the 4 m spacing used for the source designs throughout this paper and the experiments of Part II. A 4 m spacing (at 5 mCi) offers $0.4 < P_{\text{FN}} < 0.5$ for altitudes $z > 5$ m, and was deemed a reasonable tradeoff

between the accuracy of 3 m spacing and the spatial coverage of 5 m spacing.




Fig. 11. Top: P_{FN} vs array source spacing at various altitudes in the NG-LAMP hot/coldspot example of Fig. 10, with a baseline activity of 5 mCi. Bottom: as above, but with a baseline activity of 1 mCi.

In Table I we give the false negative probability P_{FN} for each detector system and source using the 4 m array source spacing (except for 2 m in the plume source) and 5 m altitude of Fig. 10. The P_{FN} were computed via the full shifted Gamma distributions discussed in Section II-D, both with and without summing the λ arrays over the J detector crystals. We exclude the un-summed MiniPRISM calculations, however, as in this case many of the counts per time bin are ~ 1 , and the distribution of the deviance becomes difficult to estimate. The plume source, which was designed to have regions of denser-than-standard (2 m) spacing, is unsurprisingly the most difficult to distinguish from its continuous source, with $P_{\text{FN}} \geq 0.49$. Of the remaining seven source patterns with only the standard 4 m spacing, the four sources with uniform activity per point—the 10×10 square, 8 and 12 m separations, and L-shape—have the next-highest P_{FN} of 0.47–0.48 for NG-LAMP (un-summed). The linear gradient and hot/coldspot then have the next-highest $P_{\text{FN}} \simeq 0.45$, and the hot line has the lowest $P_{\text{FN}} = 0.420$ (un-summed). The hot line is the most distinguishable from its continuous source likely due to the high number of passes over the areas of greatest activity (even compared to the hot/coldspot and linear gradient sources) and high count rates are where the system will have the best statistics with which

to distinguish small differences in activity concentrations.

TABLE I
FALSE NEGATIVE PROBABILITY BY SOURCE AT 5 M AGL, 4 M SPACING

	NG-LAMP	NG-LAMP (sum)	MiniPRISM (sum)
plume	0.494	0.490	0.497
10 × 10 square	0.480	0.466	0.483
8 m separation	0.476	0.461	0.479
L-shape	0.474	0.456	0.477
12 m separation	0.471	0.451	0.474
linear gradient	0.458	0.430	0.462
hot/coldspot	0.452	0.418	0.453
hot line	0.420	0.362	0.420

The summed NG-LAMP P_{FN} values are systematically lower than their un-summed counterparts, i.e., summing over detectors results in higher sensitivity to array vs continuous sources when using the deviance metric. This behavior can be understood via Section II-D—we expect that the deviance for data distributed in I bins is asymptotically distributed as a χ^2 with I degrees of freedom, and therefore summing over detectors results in fewer bins, lower variance in the deviance statistic, and thus higher statistical power. Finally, we note that the summed NG-LAMP P_{FN} are also systematically smaller, i.e., better at distinguishing array vs continuous sources, compared to the summed MiniPRISM values, because the NG-LAMP effective area to 511 keV photons is $\sim 2\times$ larger than that of MiniPRISM, and therefore NG-LAMP has higher statistical power. We expect however that when quantitative image reconstructions are performed in future work, the improved angular resolution of MiniPRISM (with un-summed data) will result in better reconstructions than NG-LAMP.

V. DISCUSSION

We have demonstrated design principles for emulating truly continuous distributed radiological sources with surrogates constructed of point source arrays. Although our point source array technique is straightforward in principle—coarse-graining or downsampling a continuous source at a fixed point density—the Poisson deviance test provides a level of statistical rigor to the question of whether the point source array is dense enough to effectively mimic the continuous source. A user can therefore construct a surrogate distributed source to a desired fidelity and maintain the safety and reconfigurability of using multiple sealed point sources.

Although the example calculations in Section IV show high fidelity to the continuous sources—namely, typical false negative probabilities of $P_{\text{FN}} > 0.45$, almost indistinguishable from chance—we emphasize that these results will only hold for the particular detector, trajectory, and source pattern under consideration. For smaller source-to-detector standoffs, e.g., in handheld or ground-vehicle-based rather than UAS-based surveys, P_{FN} will be generally much lower for the same source spacing (see Fig. 11); put another way, handheld or ground-vehicle-based surveys will require a much denser array source to maintain a high P_{FN} . We also note that our results are limited to singles-mode events, and that analyses using

Compton-mode events may require still-denser array spacings to compensate for the increased directional information available.

We have also shown via Fig. 8 that our method can be extended via Monte Carlo simulation to include experimentally-relevant but non-Poisson sources of noise such as position and activity errors in the deployed array source. Moreover, the additional deviance introduced by even $\pm 10\%$ (standard deviation) activity variation about the mean can induce substantial excess deviance and in some sense artificially decrease the P_{FN} . If such experimental errors cannot be adequately controlled in the field (or carefully modelled in simulation), then our deviance-based methods may be of only limited utility for the “inverse problem” of determining whether a measured dataset arose from a continuous or array source. Our methods are still however useful for the “forward problem” of designing surrogate array sources.

Finally, although we have focused on 2D planar arrays, our technique is not in principle limited to flat planes. Fully 3D source grids could be constructed and tested, though the additional numbers of sources and the requirement to place sources at heights up to ~ 100 m would present operational challenges. More achievable are so-called “2.5D” scenarios, which could be demonstrated by placing sources on hilly terrain and/or on collections of surfaces such as vehicles and buildings. Such 2.5D scenarios would be particularly interesting for emergency-response-relevant tests and demonstrations.

ACKNOWLEDGEMENTS

This material is based upon work supported by the Defense Threat Reduction Agency under HDTRA 13081-36239. This support does not constitute an express or implied endorsement on the part of the United States Government. Distribution A: approved for public release, distribution is unlimited.

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

This research used the Lawrence computational cluster resource provided by the IT Division at the Lawrence Berkeley National Laboratory (Supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231).

REFERENCES

- [1] Blake Beckman, Anna Rae Green, Laurel Sinclair, Blaine Fairbrother, Tim Munsie, and Dan White. Robotic dispersal technique for 35 GBq of ^{140}La in an L-polygon pattern. *Health physics*, 118(4):448–457, 2020.
- [2] Heather M Pennington, Stephanie Neuscammann, Craig Tenney, Jacqueline Brandon, Nick Mann, and John Giles. Microscale radiological dispersal model validation. Technical report, Sandia National Laboratories, 2020.
- [3] Anna Rae Green, Lorne Erhardt, Luke Lebel, M John M Duke, Trevor Jones, Dan White, and Debora Quayle. Overview of the full-scale radiological dispersal device field trials. *Health physics*, 110(5):403–417, 2016.
- [4] US Department of Homeland Security. Radiological dispersal device (RDD) response guidance: Planning for the first 100 minutes, 2017. Retrieved Oct. 5, 2021, from https://www.dhs.gov/sites/default/files/publications/NUSTL_RDD-ResponsePlanningGuidance-Public_171215-508.pdf.
- [5] Mathew W Swinney, Douglas E Peplow, Bruce W Patton, Andrew D Nicholson, Daniel E Archer, and Michael J Willis. A methodology for determining the concentration of naturally occurring radioactive materials in an urban environment. *Nuclear Technology*, 203(3):325–335, 2018.
- [6] Nathanael Simerl, Jace Beavers, Jacob Milburn, Miranda Dodson, Ryan Strahler, Richard Kroeger, Ivan Ulloa-Garcia, Bryan Moosman, Terence Sin, Jeffrey Kagan, et al. Contamination measurements from simultaneous activated potassium bromide radiological dispersal devices with a collimated vehicular sensor. *Health Physics*, 120(6):618–627, 2021.
- [7] R Pavlovsky, JW Cates, WJ Vanderlip, THY Joshi, A Haefner, E Suzuki, R Barnowski, V Negut, A Moran, K Vetter, et al. 3D Gamma-ray and Neutron Mapping in Real-Time with the Localization and Mapping Platform from Unmanned Aerial Systems and Man-Portable Configurations. *arXiv:1908.06114*, 2019.
- [8] R. T. Pavlovsky, J. W. Cates, M. Turqueti, D. Hellfeld, V. Negut, A. Moran, P. J. Barton, K. Vetter, and B. J. Quiter. MiniPRISM: 3D Realtime Gamma-ray Mapping from Small Unmanned Aerial Systems and Handheld Scenarios. In *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Manchester, UK, 2019.
- [9] L. A. Shepp and Y. Vardi. Maximum Likelihood Reconstruction for Emission Tomography. *IEEE Trans. on Medical Imaging*, 1(2):113–122, 1982.
- [10] D. Hellfeld et al. Gamma-ray point-source localization and sparse image reconstruction using Poisson likelihood. *IEEE Trans. Nucl. Sci.*, 66(9):2088–2099, Sept. 2019.
- [11] Jayson R Vavrek, Daniel Hellfeld, Mark S Bandstra, Victor Negut, Kathryn Meehan, William Joe Vanderlip, Joshua W Cates, Ryan Pavlovsky, Brian J Quiter, Reynold J Cooper, et al. Reconstructing the position and intensity of multiple gamma-ray point sources with a sparse parametric algorithm. *IEEE Transactions on Nuclear Science*, 67(11):2421–2430, 2020.
- [12] JH Hubbell, RL Bach, and JC Lamkin. Radiation field from a rectangular source. *J. Res. NBS C*, 64:121–138, 1960.
- [13] AE Proctor. Aerial radiological surveys. Technical report, Bechtel Nevada Corp., NV (US); Nevada Test Site, Mercury, NV (US), 1997.
- [14] John A Auxier. *Experimental evaluation of the radiation protection afforded by residential structures against distributed sources*. Oak Ridge National Laboratory, 1959.
- [15] Tenzing Joshi, Brian Quiter, Joey Curtis, Mark Bandstra, Reynold Cooper, Daniel Hellfeld, Marco Salathe, Alex Moran, Jayson Vavrek, Department of Homeland Security, DOD Defense Threat Reduction Agency, and USDOE. Multi-modal Free-moving Data Fusion (MFDF) v1.0, 2020.
- [16] S Agostinelli, J Allison, K Amako, J Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G Barrand, et al. Geant4—a simulation toolkit. *Nucl. Instrum. Methods Phys. Res. A*, 506(3):250–303, 2003.
- [17] J Allison, K Amako, J Apostolakis, H Araujo, PA Dubois, M Asai, G Barrand, R Capra, S Chauvie, R Chytrcek, et al. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.
- [18] J Allison, K Amako, J Apostolakis, P Arce, M Asai, T Aso, E Bagli, A Bagulya, S Banerjee, GJN Barrand, et al. Recent developments in Geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.
- [19] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2D LIDAR SLAM. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278. IEEE, 2016.
- [20] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part I. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.
- [21] Tim Bailey and Hugh Durrant-Whyte. Simultaneous localization and mapping (SLAM): Part II. *IEEE robotics & automation magazine*, 13(3):108–117, 2006.
- [22] THY Joshi, D Hellfeld, MS Bandstra, WJ Vanderlip, K Meehan, JW Cates, RT Pavlovsky, BJ Quiter, RJ Cooper, and K Vetter. Experimental demonstration of additive point source localization. In *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Manchester, UK, 2019.
- [23] Daniel Hellfeld, Mark S Bandstra, Jayson R Vavrek, Donald L Gunter, Joseph C Curtis, Marco Salathe, Ryan Pavlovsky, Victor Negut, Paul J Barton, Joshua W Cates, et al. Free-moving quantitative gamma-ray imaging. *Scientific reports*, 11(1):1–14, 2021.
- [24] Steve Baker and Robert D. Cousins. Clarification of the use of CHI-square and likelihood functions in fits to histograms. *Nuclear Instruments and Methods in Physics Research*, 221(2):437–442, April 1984.
- [25] P. McCullagh and John A. Nelder. *Generalized Linear Models, Second Edition*. CRC Press, August 1989.
- [26] R. A. Fisher. The significance of deviations from expectation in a Poisson series. *Biometrics*, 6(1):17–24, 1950.
- [27] Carolyn M Chen, Laurel E Sinclair, Richard Fortin, Maurice Coyle, and Claire Samson. In-flight performance of the Advanced Radiation Detector for UAV Operations (ARDUO). *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 954:161609, 2020.