Lawrence Berkeley National Laboratory

Title

Surrogate distributed radiological sources II: aerial measurement campaign

Permalink

https://escholarship.org/uc/item/0xd613c0

Authors

Vavrek, Jayson R
Corey Hines, C
Bandstra, Mark S
et al.

Publication Date

2024

DOI

10.1109/tns.2024.3351611

Copyright Information

Peer reviewed

Surrogate distributed radiological sources II: aerial measurement campaign

Jayson R. Vavrek, C. Corey Hines, Mark S. Bandstra, Daniel Hellfeld, Maddison A. Heine, Zachariah M. Heiden, Nick R. Mann, Brian J. Quiter, and Tenzing H.Y. Joshi

Abstract—In this second part of a multi-paper series, we present results from outdoor aerial measurements of surrogate distributed gamma-ray sources. We detail the design, manufacture, and testing of 300 individual ∼7 mCi Cu-64 sealed sources at the Washington State University research reactor, and their deployment in various source patterns (each comprising up to 100 point sources) during the aerial measurement campaign. We show the results of two such measurements, in which approximate source shapes and qualitative source intensities can be seen from the count rate vs. position plots, even without performing reconstructions. We also detail our efforts in ground-truthing the deployed sources and comparing measured gamma ray data to model predictions. In particular we compare measured vs. expected count data using the Poisson deviance formalism of Part I to evaluate whether the fielded surrogate point-source arrays "look like" their truly continuous distributed source analogues. More generally, we find that the point-source array technique provides high source placement accuracy, relative ease of quantifying the true source configuration, scalability to source dimensions of ≲100 m, ease of reconfiguration and removal, and relatively low dose to personnel. Finally, we consider potential improvements and generalizations of the point-source array technique for future measurement campaigns.

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

I. INTRODUCTION

AS described in Part I of this multi-paper series, an important capability in radiological emergency response is to quantitatively map distributed sources of radiation. Creating such distributed sources for developing and testing detectors and reconstruction algorithms is difficult, however, due to the numerous challenges posed by truly continuous distributed sources. In Part I we described how to build a surrogate distributed source from an array of point sources, and, more importantly, test whether the surrogate array source "looks like" its continuous source analogue in a statistically-founded way. In this Part II, we describe an aerial measurement campaign at Washington State University (WSU) in which we measured several different surrogate array sources each comprising up to 100 ∼7 mCi (∼259 MBq) Cu-64 sources. We detail our efforts in producing the sources, comparing measured absolute gamma count rates against model predictions, and applying the Poisson deviance framework of Part I to evaluate whether the surrogate point source array "looks

like" a truly continuous distributed source. We then conclude with a brief discussion of systematics and recommendations for future surrogate distributed source experiments.

II. SOURCE FABRICATION AND ASSAY

A. Source fabrication

Cu-64 was chosen as an attractive nuclide for the distributed sources measurement campaign for several reasons. First, the β^+ decay of Cu-64 to Ni-64 results in a prominent annihilation photon line at 511 keV with a yield of 0.352 photons per disintegration [1], [2]. This energy is suitable for both singles and Compton imaging, and is close to the 662 keV line from long-lived Cs-137 contamination following reactor accidents such as those at Chernobyl and Fukushima [3]. Second, the weak 1346 keV line provides a convenient check on the analysis, albeit at very limited statistics, and does not strongly interfere with measurements of the 511 keV photons. Third, high-elemental-purity (99.999–99.9999%) copper pellets (natural abundances 69.15% Cu-63, 30.85% Cu-65) were commercially obtainable from American Elements [4], reducing impurities produced during the neutron irradiation of the pellets. These high-purity pellets could also be massed to achieve one pellet per source in order to reduce handling requirements and simplify source tracking. Solid metal pellets also reduced the risk of accidental radionuclide release that would be present with powdered or liquid sources, especially in an outdoor setting. Finally, the short half-lives of the neutron capture products Cu-64 ($t_{1/2} = 12.7$ hours) and Cu-66 ($t_{1/2} = 5.10$ minutes) ensured that relatively large total activities could be produced without any long-lived radioactive waste.

To produce the Cu-64 sources, three batches of 100 high-purity copper pellets were irradiated for up to 2400 s by the thermal and epithermal neutron flux of the WSU Nuclear Science Center 1 MW TRIGA reactor. Following irradiation, the copper pellets were stored in the reactor pool; approximately six hours after irradiation, each pellet was transferred into a 2/5-dram vial that was pre-epoxied into a 2-dram vial (see Fig. 1). Once the copper pellet was transferred, the remainders of its 2/5- and 2-dram vials were filled with epoxy. After each 100-source batch was finished and distributed into epoxied vials, the sources were left to cure overnight. The next morning the 2-dram vials containing the sources were capped, heat sealed, and checked for contamination. Contamination swipes found no removable contamination present. The sealed and cured material was loaded into tennis balls with an opening

slit cut into them for transport and use on the field. Tennis balls were chosen for their ease of both visually tracking on the field during the exercises and ease of manipulation and replacement by long-handled grabber tools. In total, the sources were allowed to cool for approximately a day between irradiation and use on the field, allowing the Cu-66 component ($t_{1/2} = 5.10$ minutes) to completely decay out.
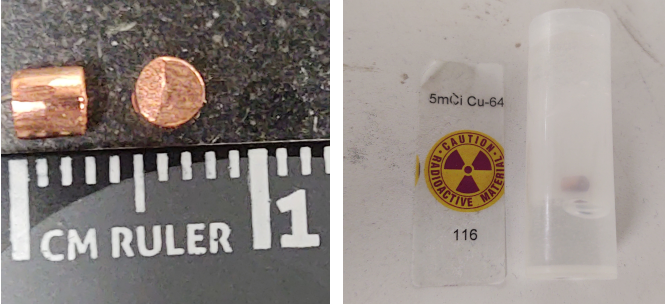

Fig. 1. Left: two representative copper pellets. Right: a copper pellet encapsulated in epoxy and double-encased in vials.

B. Source assay

Following the distributed sources measurement campaign, the Cu-64 sources were assayed via high-purity germanium (HPGe) gamma spectroscopy to determine the ground truth pellet activities. In the initial post-experiment gamma assay, however, the sources were placed too close to the HPGe detector, causing the variations in the copper pellet position within the epoxy embedding to induce large, uncontrolled variations in the source-to-detector distance compared to the source positions of the available calibration standards. These distance variations caused the assays to measure 50% relative activity variations across samples that were irradiated in very similar conditions and were expected to have very nearly uniform activity. As a result, 18 representative copper pellets were re-irradiated under reactor conditions as similar as possible to the first irradiations, with minimal expected changes to the neutron flux due to fuel burn-up in the interim. After cooling for 165 hours, the 18 pellets—this time not encased in epoxy—were individually surveyed by three independent HPGe detectors, each with its own dedicated efficiency calibration. During the 300 s assays, $\sim 2 \times 10^5$ net counts were recorded in the 511 keV peak per pellet, with dead times of $\sim 5\%$.

The activities assayed in this second (henceforth "September") batch of pellets were then used to determine the expected activities produced by the first (henceforth "August") set of irradiations. In particular, the irradiation process can be described by a neutron point-kinetics model for the Cu-64 population in a copper pellet as a function of time. We define $\phi(t)$ as the one-group mean apparent neutron flux across the volume of the copper pellet, rather than the true incident flux on the boundary of the copper cylinder. As such, we do not need to explicitly account for correction factors such as self-shielding [5]. The Cu-64 number density $N_{64}(t)$ during irradiation is then governed by the differential equation

$$\frac{\mathrm{d}N_{64}(t)}{\mathrm{d}t} = \phi(t)N_{63}(t)\sigma_C - \lambda_{64}N_{64}(t), \qquad (1)$$

where N_{63} is the Cu-63 number density, $\sigma_C \simeq 4.5$ b [6] is the (thermal group) Cu-63 $\rightarrow$ Cu-64 radiative capture cross section, and λ_{64} is the Cu-64 decay constant. We note that the neutron-induced destruction of Cu-64 is negligible. Extending Eq. 1 to include such a term effectively modifies the destruction coefficient $\lambda_{64} \rightarrow \lambda_{64} + \sigma_D\phi$, where σ_D is the Cu-64 destruction cross section. The full σ_D is unknown, but the radiative capture component has been measured to be 270 ± 170 b [7]. Given a thermal flux of $\sim 5 \times 10^{16}$ neutrons/m^2/s, the resulting change in destruction rate is $\sim 0.01\%$.

The activity $A_{64}(t)$ associated with the number density of Eq. 1 is

$$A_{64}(t) = N_{64}(t)\lambda_{64}m_{\text{pellet}}/\rho_{\text{Cu}}, \qquad (2)$$

where m_{pellet} is the mass of the given copper pellet and ρ_{Cu} is the density of copper. Assuming that the flux is constant in time $\phi(t) \equiv \phi$, that the change in the Cu-63 population is negligible $N_{63}(t) \equiv N_{63}$, and that the initial Cu-64 population is $N_{64}(0) = 0$, we have

$$N_{64}(t) = \frac{\phi N_{63}\sigma_C}{\lambda_{64}}\left[1 - e^{-\lambda_{64}t}\right]. \qquad (3)$$

We note that Cu-63 term N_{63} can be further expanded in terms of the Cu-63 natural abundance ratio f_{63}, the molar mass of copper A_{Cu}, and Avogadro's number N_{Av} as

$$N_{63} = f_{63}\rho_{\text{Cu}}N_{\text{Av}}/A_{\text{Cu}}. \qquad (4)$$

Given irradiation ("cook") and decay ("cool") time durations Δt_{cook} and Δt_{cool}, respectively, we can write

$$N_{64}(t_f) = \frac{\phi N_{63}\sigma_C}{\lambda_{64}}\left[1 - e^{-\lambda_{64}\Delta t_{\text{cook}}}\right]e^{-\lambda_{64}\Delta t_{\text{cool}}}. \qquad (5)$$

where t_f denotes the "final" or "field" time at the end of the cooling window and thus the start of the September HPGe assay or August UAS measurement day. We can then rearrange Eq. 5 for the apparent neutron flux

$$\phi = \frac{N_{64}(t_f)}{N_{63}}\frac{\lambda_{64}}{\sigma_C}\left[1 - e^{-\lambda_{64}\Delta t_{\text{cook}}}\right]^{-1}\left[e^{-\lambda_{64}\Delta t_{\text{cool}}}\right]^{-1}. \qquad (6)$$

Eqs. 5 and 6 thus form a pair of forward and backward models that can be used to determine the August pellet activities based on the September assay data. In Eq. 6, the activity $A_{64}(t_f)$—and thus number density $N_{64}(t_f)$ via Eq. 2—is determined from the (net) number of Cu-64 counts C_{64} observed during an HPGe spectroscopy measurement, using either the 511 keV or 1346 keV spectral line. To first order (i.e., assuming the assay livetime Δt_{assay} is short compared to the Cu-64 lifetime), we can write

$$C_{64} = A_{64}(t_f)\Delta t_{\text{assay}}\epsilon b \qquad (7)$$
$$= N_{64}(t_f)V_{\text{pellet}}\lambda_{64}\Delta t_{\text{assay}}\epsilon b \qquad (8)$$

where ϵ is the total detection efficiency (at either 511 keV or 1346 keV, depending on which region of interest is used to define the counts C_{64}), and b is the branching ratio or gammas per decay again depending on the spectral line assayed. We note that while the approximate Eq. 7 is shown for clarity of presentation, our analyses correct for the slight ($\sim 1\%$) decay of the source activity during the measurement. The net

number of counts C_{64} in each measurement is determined by integrating either the 511 keV or 1346 keV peak region of interest (ROI) and subtracting a small constant background based on the ∼5 keV windows on either side of the ROI. As discussed further in Section VI, the 511 keV annihilation peak is broader than the HPGe resolution for a nuclear decay line at the same energy, and thus requires a broader ROI to accurately determine the net counts. Once the C_{64} are determined, the mean apparent neutron flux ϕ can then be computed from the September irradiations via Eq. 6, and substituted back into Eq. 5 with the August Δt_{cook} and Δt_{cool} to determine the expected August pellet number densities $N_{64}(t)$ and thus activities at the start of each UAS measurement day in Table I.

Mean apparent fluxes ϕ determined via Eq. 6 for each September pellet are shown in Fig. 2. Averaging over pellets, the mean and standard deviation of fluxes $\langle\phi\rangle$ computed from the 1346 keV line are $(5.407\pm0.441)\times10^{16}$ neutrons/m^2/s, or $(5.299\pm0.418)\times10^{16}$ neutrons/m^2/s from the 511 keV line. We note that these standard deviations arise from the variations across pellets, and not the counting statistics uncertainty shown for each point in Fig. 2. The reason for lower fluxes computed using one of the three HPGe detectors (detector 1) is unknown, but included as a systematic uncertainty in Section VI. The resulting average activity values for the start (0800 PDT) of each flight day, as computed using the 511 keV line, are shown in Table I. Average activity values computed with the 1346 keV line are consistent with those from the 511 keV line to within ∼2%.
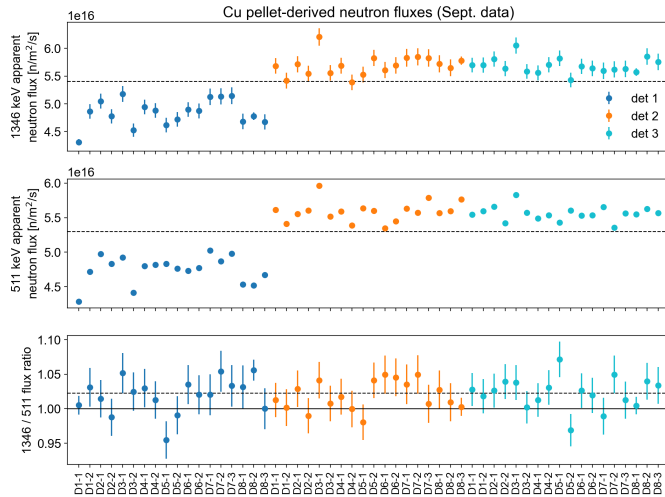

Fig. 2. Apparent fluxes derived for each of the 18 individual September pellets, using three different HPGe detectors, using either the 1346 keV peak (top) or the 511 keV peak of Cu-64 (middle). Dashed lines show the averages of the mean apparent fluxes for each energy, 5.407×10^{16} neutrons/m^2/s and 5.299×10^{16} neutrons/m^2/s for 1346 and 511 keV, respectively. The 1346/511 ratio is also shown (bottom), and is systematically slightly higher than 1.

III. SOURCE DEPLOYMENT

The WSU measurement campaign was conducted at an outdoor rugby field (GPS $46.7346°$, $-117.1474°$) near the WSU Nuclear Science Center from August 8 to 13, 2021, with source measurements on August 9, 11, and 13. Source

TABLE I
SUMMARY OF COPPER PELLET ACTIVITIES

date in Aug. '21	src configs deployed	pellet mass (avg) [g]	pellet mass (std dev) [%]	pellet activity* (avg) [mCi]
9	square L-shape	0.2173	2.43	8.466
11	8 m separation 12 m separation plume	0.1746	2.00	6.721
13	hot/coldspot linear gradient hot line	0.1752	3.37	6.810

* at 0800 PDT of each experiment day.

positions on the field were set out with marking flags prior to each source measurement day. The bottom left corner $(x,y)=(65,10)$ m common to all sources was measured via tape measure from the southwest corner of the field. The bottom right corner was marked out in a similar fashion, after which three tape measures were used to triangulate further source boundaries. The remainder of the source flags were then placed at 4 m intervals between boundaries, using a tape measure to check distance and a taut string to check linearity. This process was used to first place flags for the 10×10 source, which contains many of the source positions of the remaining seven sources, and then repeated as necessary to extend the grid for the L-shape, hot line, and rectangle separation sources. A similar process was used to place the 2 m-spaced flags for the plume source. Based on cross-checks of the diagonal distances, we estimate the flags were placed with a maximum error of ∼4 cm over ∼40 m, with most flags accurate to less than half that error. Retroreflective vinyl and traffic cones were also placed at the corners of each source distribution to help localize the source extent in the measured LiDAR point clouds. The LiDAR units measure the strength of their 905 nm laser returns, thus including these strongly-returning retroreflective materials provides more spatial context for the measurements.

On source measurement days, the sources were deployed by a team of ∼10 personnel equipped with long grabber tools to keep the radiation dose to any one person as low as reasonably achievable. Source tennis balls were placed flush with the stem of the flag when possible, but their alignment with the grid was generally not consistent. As a result, each true source position lies at approximately one tennis ball radius (measured to be ∼3.3 cm, consistent with the 3.27–3.43 cm specified by the International Tennis Federation [8]) from the flag stem. The tennis balls are, however, visible in aerial photogrammetry—see Fig. 3—allowing for small position corrections if necessary. Given the 4 m source separation and the goal of reconstructing distributed sources, such uncertainty is expected to be negligible. As described in greater detail in Part I, we then flew the NG-LAMP (4 large CLLBC modules) [9] and MiniPRISM (58 small CZT modules) [10] radiation mapping systems over the source distributions at altitudes of 5-15 m, raster pitches of 5.2 m, and speeds of 2.6 m/s.
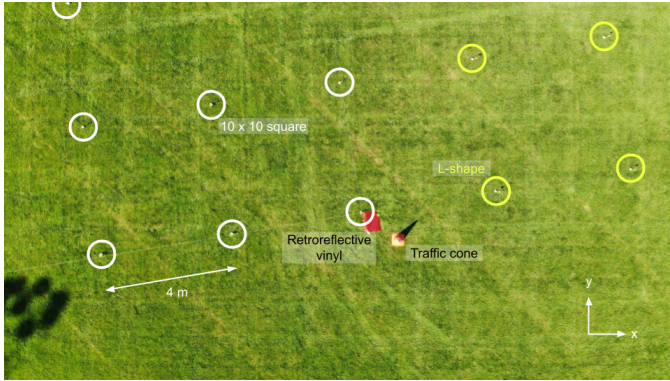
Fig. 3. Annotated aerial photograph during the deployment of the L-shape source. The white circles mark the source positions at the bottom right corner of the 10×10 square source that are also part of the L-shape, while the yellow circles mark the positions that are only part of the L-shape. Tennis balls containing sources are visible within the circles upon close inspection. The retroreflective red vinyl and orange traffic cone are also annotated. The dark blotches in the lower left are shadows of nearby floodlight structures.

IV. POST-PROCESSING

Several post-processing steps were performed on the data in order to account for non-ideal detector performance.

Both the NG-LAMP and MiniPRISM energy spectra originally exhibited gain shifts that varied across individual detector elements, broadening the summed photopeak and shifting it away from 511 keV. To compensate, we find the per-element linear gain shift necessary to return each photopeak to 511 keV, and apply that to the detector's energy data. To avoid aliasing the energy data in the process, we also apply a zero-mean Gaussian blur of 1 keV standard deviation to the listmode energy data, which we note is much less than the expected photopeak standard deviations of $\mathcal{O}(10 \text{ keV})$. Similarly, for NG-LAMP, a small time noise was applied to plots of the ROI counts vs. measurement time to avoid time bin aliasing.

Some NG-LAMP and MiniPRISM detector elements were suffering from large amounts of electronic noise or were not reading out data during several of the UAS experiments. In these cases, data from the problematic detector elements and the corresponding contributions to computed sensitivity are excluded from the analysis.

Similarly, in the MiniPRISM experiment shown later in Section V, occasional ~ 0.1 s drops in output data were observed, and appear to be due to data throughput issues on the system's on-board computer. Before running the deviance analyses of Section V-B, we identify these data drops via a median filter and exclude data within a window of ± 1 s surrounding the drop.

In addition, we have observed in prior tests that MiniPRISM can exhibit a time delay between the radiation timestamps and trajectory timestamps that needs to be calibrated for. We compute the optimum delay for each run by minimizing the deviance between the expected and shifted observed counts via `lmfit` [11], and find that the radiation data typically lags the trajectory data by ~ 0.5 s. In the later analysis of constant scale factors between expected and measured count data, the time shift is optimized simultaneously with the scale factor.

Although radiation measurements were collected during the entire UAS flight, the takeoff and landing segments were cut from the analyses of Section V in order to focus on the constant-altitude raster pattern measurements.

Coordinate transforms between the measured data and the idealized field coordinate system (see Part I) used throughout this work were determined in two steps. First, we aligned a UAS trajectory reconstructed by LiDAR simultaneous localization and mapping (SLAM) to the same trajectory as measured by the real-time kinetic (RTK) positioning system to obtain RTK $\leftrightarrow$ LiDAR transforms for each day. The optimal alignment was obtained by minimizing the sum of squared distances between the two sets of trajectory coordinates (interpolated to the same timestamps) over four parameters, the xyz translation and yaw. The LiDAR $\leftrightarrow$ field transforms were then estimated by picking points in the LiDAR point clouds at the western corners of the field boundary and aligning their vector with the $+y$ axis of the field coordinate system. This two-step process was required since there are no clear features in the RTK trajectories that reliably map onto the field coordinate system in order to directly build an RTK $\leftrightarrow$ field transform. Using the SLAM map on its own without the RTK data was challenging due to occasional trajectory reconstruction errors likely resulting from the lack of point cloud features in much of the field. Tuning the SLAM reconstruction therefore remains an important avenue in our future work.

V. RESULTS

A. Measurement results

Fig. 4 shows the results of an 8 m above ground level (AGL) MiniPRISM raster over the hot/coldspot source shown in Part I. Counts in the 511 keV photopeak ROI are plotted vs. position and vs. time with a time binning of $\Delta t = 0.2$ s; even without performing a MAP-EM reconstruction, this method of visualization produces the overall rotated square shape and hotspot. Without the reconstruction, however, the coldspot is not easily discernable from the rest of the square, and the edges are quite blurred. Moreover, this visualization is not quantitative in terms of source activity. Strong modulations in measured ROI counts are however visible in the count rate vs. time plot, indicating good contrast between regions near and far from the source distribution. Expected count rates shown in Fig. 4 are computed using the methods of Part I, forward projecting the array source to the measured trajectory, additionally accounting for decay corrections from the initial activities in Table I as per Ref. [12, Appendix C]. The measured and expected count rates across the measurement agree to 36%, as determined by assuming there is a constant scale factor between the two background-subtracted rate vs. time curves, and computing the optimum scale factor that minimizes the Poisson deviance between the scaled expected counts and the data counts (middle of Fig. 4). Possible reasons for the residual 36% discrepancy are discussed further in Section VI. After scaling the expected counts by the optimal scale factor, measurement and expectation agree quite well as shown in the ratio and discrepancy plots of Fig. 4—most discrepancies appear randomly distributed within $\pm 2\sigma$, and

the large negative outlier is due to one of the brief data drops described in Section IV and removed in post-processing. The small amount of remaining structure in the discrepancies, whether viewed as a ratio or in terms of σ, is approximately symmetric about the optimum scaling factor, and thus does not change our conclusion that a constant scale factor is appropriate. We attribute this structure to slight misalignments in the synthetic and measurement coordinate frames. Finally, in the measured spectrum (bottom of Fig. 4), the 511 keV photopeak has good contrast vs. background; small clusters of counts corresponding to the 1346 keV line from the Cu-64 sources and the 1461 keV line from the K-40 background are also visible.

Similarly, Fig. 5 shows the results of a 6 m (AGL) NG-LAMP raster over the L-shape source shown in Part I. Again, the overall L-shape is visible, but in this counts vs. position visualization, the shape is blurred over a much larger region than the true source distribution. Due to the larger extent and lower maximum concentration of the source, the contrast in the count rate vs. time plot (again, $\Delta t = 0.2$ s) is not as strong as in Fig. 4. The agreement between measured and expected count rates is 34%, similar to the 36% of Fig. 4, and again the scaled expected counts agree well with the measured data. In the detector spectra, the 511 keV and 1346 Cu-64 peaks are again prominent above background, and the K-40 1461 keV peak and La-138 1436 keV self-activity peak in the CLLBC are clearly visible though not separable.

B. Deviance analysis

We now apply the Poisson deviance framework of Part I to the results of Section V-A in order to examine whether the array sources "look like" their continuous source analogues. As described in Section IV, we use the post-processed data after removing poorly-performing detectors, cutting out intervals of dropped data, fitting the MiniPRISM time delay, and applying the optimum scale factor. These post-processing steps are necessary as the Poisson deviance framework assumes that Poisson noise and slight mean count rate variations between the array and continuous sources are the only sources of deviance in the measurement, which does not reflect reality. We note that in the following, we will use definitions of various deviance terms D defined in Part I, so we refer the reader to Section II-D of that paper.

Fig. 6 shows the deviance analysis for the NG-LAMP L-shape measurement of Fig. 5, using a concatenation of counts across the two functioning NG-LAMP detector elements. The unit deviances between the post-processed measured and expected counts are shown distributed in xy, as a histogram, and as an accumulation over the measurement duration. Modest amounts of excess deviance accumulate across the entire measurement at a roughly constant rate, leading to an observed total deviance that is inconsistent (by $\sim$1000) with both the central and non-central deviance distributions $D(\boldsymbol{n}_0|\boldsymbol{\lambda}_0)$ and $D(\boldsymbol{n}_1|\boldsymbol{\lambda}_0)$. The measured data is therefore inconsistent with both the continuous and array sources under the strict assumption that only Poisson noise contributes to the deviance. Again, because there are experimental realities beyond Poisson
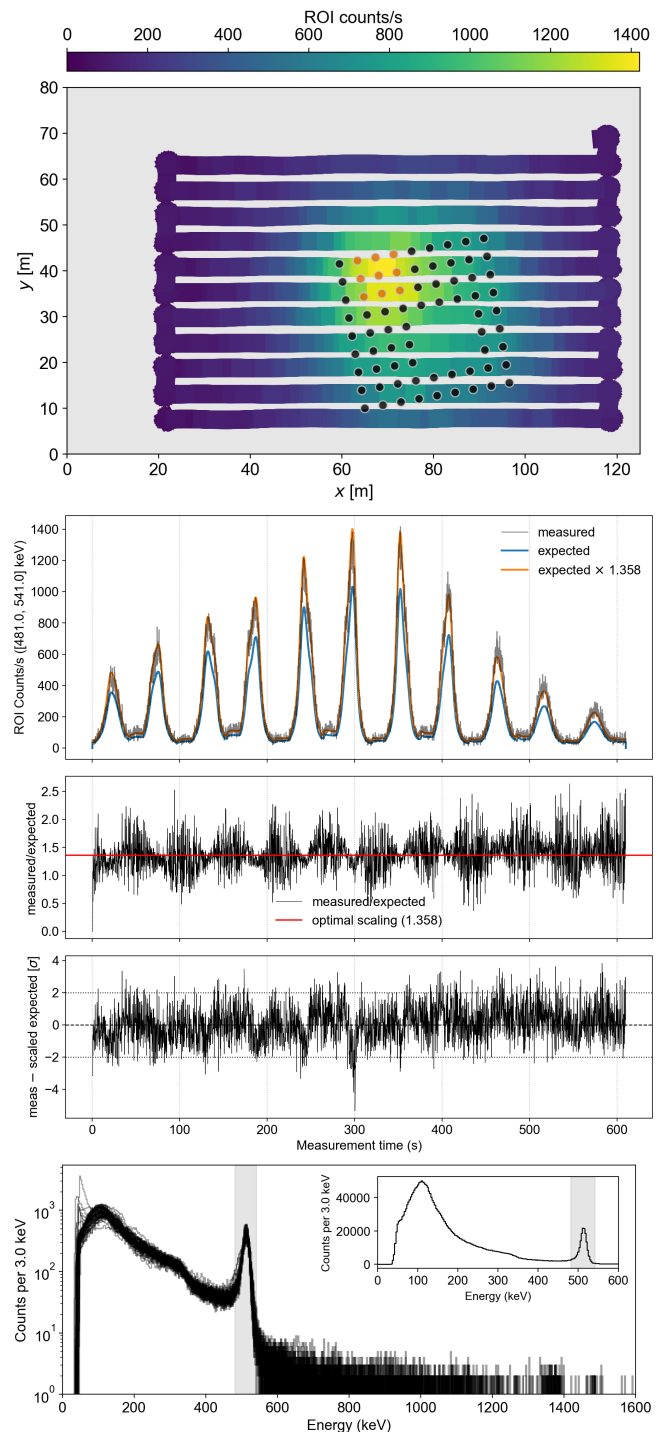

Fig. 4. Top: top-down view of a MiniPRISM measurement 8 m AGL over the hot/coldspot source. The thick curve shows the MiniPRISM detector trajectory (colorized by the measured ROI count rate at each 0.2 s pose), over the (synthetic) source (black and orange points, as also shown in Part I). Top middle: Count rates vs. time, summed over detector elements. The black curve shows the measured counts in the 511 keV photopeak ROI, while the blue curve shows the expected counts computed by forward projecting the synthetic array source onto the measured trajectory. The orange curve shows the expected counts after scaling the net signal rate by a constant scale factor of 1.358, which was found to most closely match the data. Middle: Ratio of measured to expected counts, and the optimal scaling factor between them. Bottom middle: Discrepancy (in σ) between the measured data and the scaled expected values. Bottom: energy spectra, separated by detector elements and shown on a log scale. The shaded region shows the 511 keV photopeak ROI. Bottom inset: energy spectrum, summed over detector elements and shown on a linear scale.
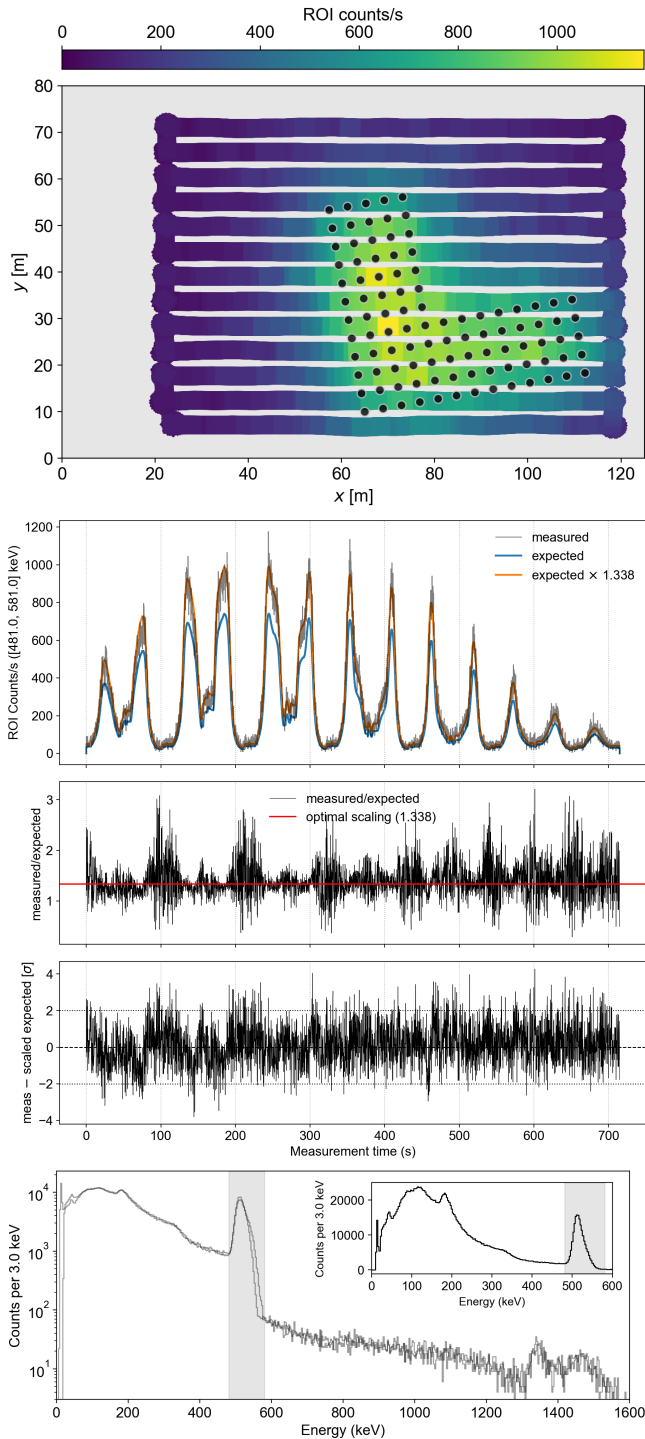
noise (e.g., coordinate frame misalignments), even after post-processing the data, this inconsistency is unsurprising.

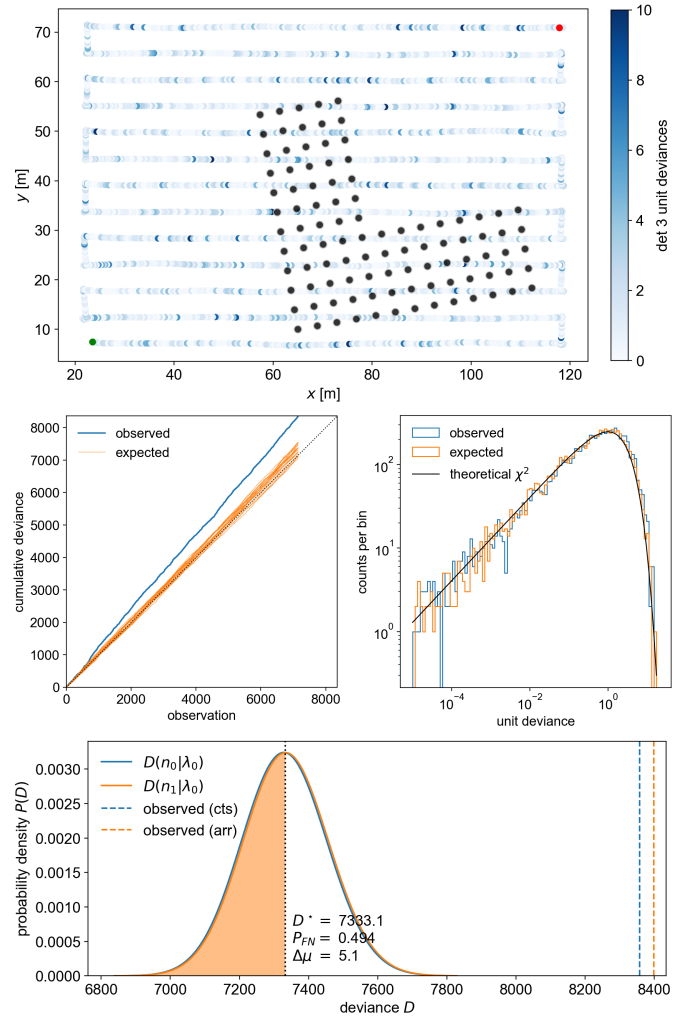Fig. 5. As Fig. 4, but for an NG-LAMP measurement of the L-shape configuration at a height of 6 m AGL.

Fig. 6. Deviance analysis for the NG-LAMP L-shape example of Fig. 5. Top: unit deviances observed at each xy position in one NG-LAMP detector. Left middle: cumulative deviances when concatenating the counts across the two NG-LAMP detectors. The blue curve shows the observed cumulative deviance while the orange curves show 20 Poisson samples of the mean count array $\boldsymbol{\lambda}_0$ expected from the optimally-scaled continuous source. Right middle: histograms of observed unit deviances (blue) and expected unit deviances (orange) based on a single Poisson sample of $\boldsymbol{\lambda}_0$. The black curve is the theoretical χ_1^2 distribution normalized to the number of observations. Bottom: expected deviance distributions (blue and orange curves) and observed deviances (blue and orange dashed lines).

Fig. 7 shows a similar deviance analysis for the MiniPRISM hot/coldspot example of Fig. 4. We note that due to some not-yet-understood deviance behaviors at low counts, here we compute the deviances using count data summed instead of concatenated across all detectors. In particular, when using concatenated data, the unit deviance histogram (right middle of Fig. 7) exhibits an additional pronounced "shoulder" near a unit deviance of 1. Moreover, the mean of the non-central distribution $D(\boldsymbol{n}_1|\boldsymbol{\lambda}_0)$ will occasionally be lower than that of the central distribution $D(\boldsymbol{n}_0|\boldsymbol{\lambda}_0)$ (giving false negative probabilities $P_{\mathrm{FN}} > 0.5$), which does not seem physically plausible. We expect these effects are related to low counts because (1) they do not occur in the NG-LAMP data, which

has larger crystals and thus higher count rates; (2) plotting the unit deviances vs. time suggests the "shoulder" of near-1 unit deviances stems from count rate valleys; and (3) performing the same analysis with $\Delta t = 1$ s time bins instead of 0.2 s substantially improves (but does not completely eliminate) the shoulder. Similar to the NG-LAMP deviance analysis, when using the summed MiniPRISM data, the excess deviance accumulates at a very roughly constant rate, though moderate rate increases over the hottest parts of the source are noticeable, and the final observed deviance is inconsistent with both the continuous and array sources under the Poisson-noise-only assumption.
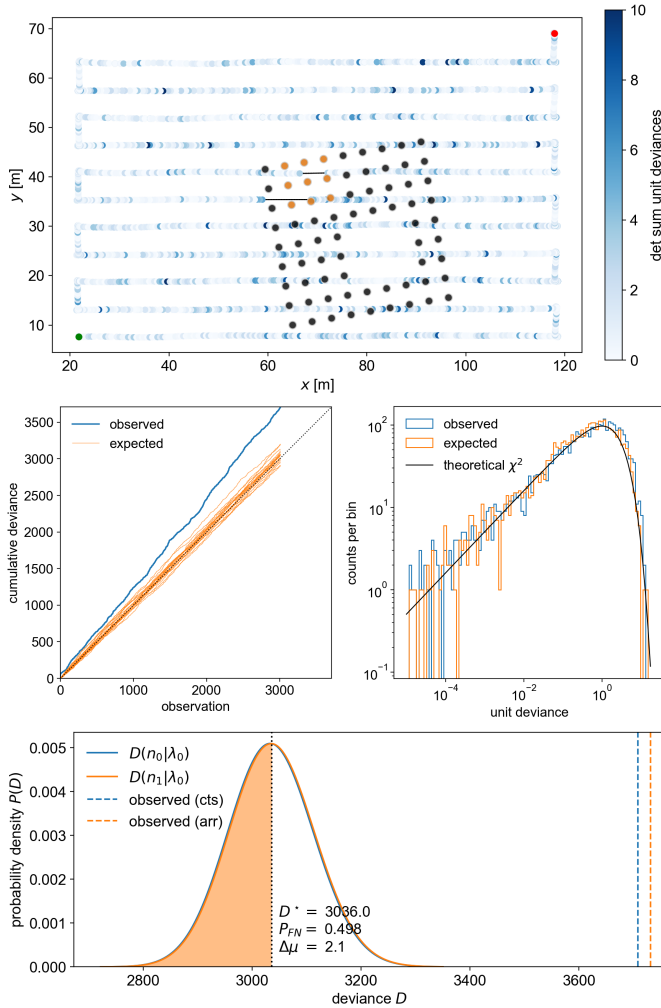
Fig. 7. As Fig. 6, but for the MiniPRISM measurement in Fig. 4 and using count data summed instead of concatenated over detectors. The gaps in the trajectory colorization indicate where the data was cut before use in the deviance analysis.

VI. DISCUSSION

A. Additional corrections and uncertainties

The results of Section V-A—in particular, Figs. 4 and 5—show $\sim$35% agreement between the measured and expected 511 keV photopeak count rate. Agreement to this level is in fact seen across the entire set of measurements (in which reliable RTK position data is available), and is relatively constant across measurement days, source configurations, and detector systems. In particular, the average levels of agreement for the NG-LAMP, MiniPRISM, and overall datasets are 35.2%, 33.7%, and 34.7%, respectively. For completeness, though, we note that we have analyzed and either corrected for or ruled out several possible sources of error, and made estimates of our dominant systematic uncertainties:

1) **attenuation in air:** although most flight altitudes were $\lesssim$10 m, the distributed nature of the source and wide raster patterns mean that source-to-detector distances are often on the order of the 511 keV mean free path in air, 96.2 m for dry air near sea level [13]. Air attenuation therefore causes non-negligible reductions in the expected photopeak counts. To model air attenuation, we use the NIST XCOM mass attenuation coefficient tables for dry sea-level air [14], but choose an air density that reflects local weather conditions in Pullman, WA on the measurement days. We use an air temperature of 28 °C, pressure of 928 hPa, and relative humidity of 30% as representative values across and within all three days [15]. We then use the simplified air density formula of Ref. [16] to arrive at an air density of $\rho = 1.069$ kg/m^3. For the measurements shown in Figs. 4 and 5, the overall magnitude of the air attenuation correction is an approximately 17% reduction in the sum of the expected photopeak counts. As we have used a single air density value, changes in pressure and temperature throughout the week of measurements are expected to induce a $\sim$3% uncertainty in forward-projected counts.

2) **attenuation in the source holder:** the activities of Table I were computed for bare copper pellets without the source holders (epoxy encasing and tennis balls). To correct for attenuation losses in the source holders, we simulate monoenergetic photons emitted isotropically and uniformly throughout a 0.19 g copper pellet of diameter and height 3 mm. The copper volume is centered in an epoxy cylinder of diameter 1.5 cm and length 4 cm, surrounded by air, and placed inside a tennis ball volume comprising a 3.5 mm rubber shell and a 2.5 mm felt shell. Material compositions were generally taken from Ref. [17], though high-density (1.5 g/cm^3) rubber [18] was used to ensure the modelled tennis ball mass was consistent with its measured value of 57 g. Since both the on-field and HPGe assay measurements are subject to attenuation by the copper pellet itself, we define the transmission fraction as the ratio of full-energy photons escaping the tennis ball to those escaping the copper pellet, thereby quantifying losses from the epoxy, internal air, and tennis ball volumes only. Moreover, since the orientations of the epoxy cylinders were randomized during the field measurements, the transmission fraction is taken as an average over all emission directions. We find that the attenuation losses are primarily driven by the epoxy cylinder, and that the average transmission fractions are 0.861 for 511 keV photons or 0.910 for 1346 keV. We then multiply the pellet activities in Table I by 0.861 to obtain apparent source activities for the

forward projections in Figs. 4 and 5. We note however that while we have assumed the epoxy orientations were isotropic, the vials may have tended to settle closer to a horizontal orientation, thereby biasing the detected emission distribution away from the poles and reducing the epoxy attenuation while the detector was overhead. In the extreme case of all emissions occurring in the radial direction, the transmission fraction would increase to about 0.89 for 511 keV photons. The increase in transmission due to the epoxy orientation bias is therefore at most 3%, and likely much less.

3) **511 keV region of interest:** care must be taken when defining the 511 keV region of interest (ROI) in the HPGe ground truth activity assay. We found that the default 511 keV ROI provided by the Genie™ 2000 spectroscopy software [19] relied on a peak width calibration based on nuclear decay lines. However, the 511 keV annihilation line is subject to additional Doppler broadening [12, p. 441], rendering its width substantially larger than nearby decay lines. Because the same Genie™ 2000 analysis pipeline was used to empirically plan irradiation times and to assay the pellets used, this too-narrow ROI led to both a $\sim30\%$ overproduction of Cu-64 and a corresponding initial $\sim30\%$ underestimate of the Cu-64 ground truth activity. In our final analyses, this ROI is broadened to include the entire 511 keV peak.

4) **in-scatter:** although we have corrected our forward projections (Eq. 3 of Part I) for scattering losses between the sources and the detector, we have so far only considered scattering as a loss mechanism. However, due to the often-large source-to-detector standoffs and the non-zero width of the energy ROI used to define the measured 511 keV signal, there is a non-negligible solid angle in which small-angle in-scatter ("buildup", typically from air, but potentially also from the ground) can occur, increasing the ROI signal compared to the prediction of Eq. 3 of Part I. This effect manifests as the apparent step function underneath the 511 keV peak (most notably in Fig. 5) since Compton scattering always reduces the photon energy. To compensate for this in-scatter without running computationally intensive scattering simulations, we derive an approximate buildup correction from the global 511 keV peak (i.e., summed over the entire run) from the NG-LAMP run in Fig. 5. We fit the peak with an exponentially-modified Gaussian peak shape plus two different exponential backgrounds on either side of the centroid that are blurred together by convolution with the peak's resolution. We find that the area under this "source-induced background" component comprises approximately 9.2% of the ROI area. Although this exact value will depend on the source geometry, trajectory, and perhaps even the detector, we take it as a representative in-scatter fraction for our experiments. In the comparisons of Figs. 4 and 5, we therefore downsample the measured listmode data, randomly dropping 9.2% of events. We note that the higher-fidelity and only moderately more computationally intensive air scattering model outlined in Ref. [20] could be useful here, but would require

substantial work to integrate with our existing GPU forward projection code.

5) **altitude errors:** when dealing with a single point source, small source-to-detector distance errors have a quadratic effect on the expected counts λ. However, as shown in Eq. 2 of Part I, for distributed sources the dependence is generally logarithmic, and thus it would take altitude errors of several meters to explain the discrepancy between our measurements and expectation based on the analysis of the activity assays. Moreover, adjusting the detector altitude scales counts non-linearly with respect to source proximity, generally improving the agreement in only some parts of the count rate comparisons while worsening it in others. There is a small altitude uncertainty of around ±25 cm due to the non-zero slope of the field and the accuracy of the spatial transformations, corresponding to a roughly $\pm5\%$ change in expected counts.

6) **detector pitch:** the measurements with RTK trajectory data do not contain UAS pitch or roll info, but the detector angular response varies with system orientation. We find that adding a constant $10°$ pitch, representative of real pitches observed on the field, would change the summed forward projected counts by $<1\%$. In the extreme case of a constant $45°$ pitch, the change would be $\sim2\%$. For simplicity, therefore, no pitch was imputed to the RTK trajectories in the analyses presented.

7) **detector response validation:** we showed good agreement between data and our response models for the MiniPRISM detector in Ref. [21]. We performed an additional validation study with NG-LAMP and Na-22 check sources (which also emit at 511 keV), and found agreement between expected and measured efficiencies (defined in terms of the effective area—see Part I Section II.A) to within $\sim10\%$. Moreover, the average difference of 1.5% between the NG-LAMP and MiniPRISM datasets in this work is better than the fidelity to which we expect to know our detector responses. Finally, we note that our detector response models do not use 511 keV response simulations directly, but rather interpolate (in log-log space) between two nearby energies (356 and 662 keV for NG-LAMP, and 500 and 600 keV for MiniPRISM). Because the response efficiencies decrease roughly exponentially above ~200 keV, performing the energy interpolation in linear space instead of log-log space would lead to a $+15\%$ interpolation error with NG-LAMP (but only $+1\%$ with MiniPRISM).

8) **forward projection code:** we compared the GPU-based `mfdf` [22] forward projection code (Eq. 3 in Part I) against a separate simpler CPU-based Python implementation and found results consistent to $\sim1\%$ in benchmark cases. Moreover, we modelled with `mfdf` one of the distributed Na-22 511 keV sources in Ref. [21], and replicated the measured counts to within experimental error.

9) **decay correction to mean activity:** given that the ~15 minute measurement times are not entirely insignificant ($\sim1\%$) compared to the Cu-64 half-life ($t_{1/2} = 12.7$ hours), we decay-correct the source activity in each

measurement to the mean activity over the measurement time (following Ref. [12, Appendix C]) rather than simply the activity at, e.g., the start or midpoint time.

10) **indirect production of 511 keV photons:**

   a) pair production and scatter from the 1346 keV line in the detector and surrounding materials: as an extreme upper limit, even if every 1346 keV photon incident on the detector led to a 511 keV detection, the relative contribution to the 511 keV ROI would be limited to about 0.25% based on photon yield [1] and efficiency ratios.

   b) pair production and scatter from the 1346 keV line in the ground: we find via Monte Carlo simulation that for an isotropically-emitting plane source, each 1346 keV photon results in only $\sim 6 \times 10^{-3}$ photons leaving the ground in a 511 ± 30 keV energy window.

   c) pair production and scatter from the 1346 keV line in the source holder: based on our simulations of attenuation in the epoxy and tennis ball source holders, fewer than 10^{-3} annihilation photons are emitted from the tennis ball for every 1346 keV photon generated in the copper pellet.

Altogether, our dominant known sources of systematic uncertainty are as follows:

- detector response: $\sim 10\%$.
- Cu-64 source activities and associated mean apparent flux calculations: $\sim 10\%$ from the observed spread in activities across the three HPGe detectors.
- in-scatter: $\sim 10\%$ from the approximate nature of the global in-scatter correction.
- altitude uncertainties: $\sim 5\%$ from potential detector altitude errors.
- air density uncertainties: $\sim 3\%$ from using a single representative air density for all measurements.

Potential unknown sources of systematic uncertainty, such as the epistemic uncertainty in the source activity redetermination and the constant flux assumption of Section II-B, are not quantified. Adding these known systematics in quadrature we have a total systematic error of $\sim 18\%$, approximately half the observed $\sim 35\%$ excess counts observed in Figs. 4 and 5 and in the average NG-LAMP, MiniPRISM, and overall datasets. To improve these systematics in potential future measurement campaign, we recommend using a true nuclear decay line (e.g., the 662 keV line of Cs-137 rather than 511 keV) to avoid spectroscopy errors; ensuring the same set of sources that is used during the measurement is used for the source activity assay; and improving air models for both the air in-scatter and weather-dependent attenuation corrections. Finally, we note that while this $\sim 35\%$ discrepancy between measurements and modeling exists, since it appears to be fairly consistent across detectors and both across and within measurements, we can still quantitatively map distributed radiological sources, though some absolute activity scale correction or calibration may be required. Similarly, although we found discrepancies between the observed and predicted Poisson deviances in Section V-B, the Poisson deviance framework of Part I is still useful in the forward direction of designing fieldable surrogate distributed sources; it is not intended to generate exact predictions of measured deviances.

B. Future work

In general, our surrogate point-source array technique serves as a useful proof-of-concept for future measurement campaigns where well-controlled distributed sources would be desirable. We expect these techniques to be useful up to source dimensions of around 100 m, beyond which it may be preferable to develop a more scalable method at the expense of some spatial accuracy.

Having a team of ~ 10 personnel to deploy sources at the pre-set flags was instrumental in deploying, reconfiguring, and removing source distributions on the order of ten minutes, keeping dose to personnel relatively low in the process. The Berkeley personnel, who were often present on the field to place sources and pilot or spot the UAS, received an average (standard deviation) full-body dose equivalent of $220(17)$ μSv over the entire measurement campaign. We expect this dose could be further reduced in a truly open-field measurement with no surrounding fences or light poles, as the UAS team could safely increase their standoff to the measurement area.

We also note that for operational simplicity, we have only demonstrated point-source array designs and measurements on a flat 2D ground plane, and generally with fixed-altitude raster patterns above the source. Generalizations to arrays over hilly surfaces or full 3D environments are possible, but will require a map of the environment to be modeled or measured (e.g., via LiDAR SLAM). In fact when such a 3D map is available, it is possible to account for attenuation in the scene [23] that has made activity reconstruction difficult in similar measurement scenarios [24].

One remaining open question is the low-counts behavior of our Poisson deviance framework. As noted in the discussion around Fig. 7, we had to sum the MiniPRISM count data over the detector elements instead of concatenating as we did for the NG-LAMP data of Fig. 6, in order to avoid unrealistic false negative probabilities $P_{\text{FN}} > 0.5$. While we have identified a heuristic for when low-count issues may arise (the presence of the "shoulder" in the unit deviance histogram), a more rigorous study of the low-count regime should be undertaken.

We emphasize again that analyzing detected counts vs. xy position as in Figs. 4 and 5 does not provide quantitative estimates of the source total activity or distribution. To answer such questions, we require quantitative reconstruction methods. In an upcoming work, we will use regularized ML-EM (MAP-EM) to quantitatively reconstruct the source distributions and compare to known ground truth; such work could also involve other methods under study such as particle filters and genetic algorithms.

Finally, we recently conducted an additional distributed sources campaign in September 2023. The 2023 campaign built on the 2021 results presented here, deploying 100×0.1 mCi Cs-137 sources on hilly terrain in order to demonstrate more easily activity-ground-truthed quantitative distributed source mapping in non-planar geometries. Results from this campaign will be presented in future work.

VII. CONCLUSION

We have demonstrated a method for emulating distributed gamma ray sources with arrays of sealed point sources, and performed aerial measurements of Cu-64 source arrays of total activity up to $\sim$700 mCi. We found that the sealed array sources were easily reconfigurable, enabling multiple source configuration measurements per day, and we obtained quantitative agreement between modelled and measured count rates at the level of $\sim$35%, around double our expected systematic uncertainties. This agreement could be improved in potential future measurement campaigns by enhancing the precision and accuracy of the ground truth activity assay and by better understanding second-order effects such as air in-scatter. These measurements will form the basis of upcoming studies applying quantitative reconstruction techniques to accurately determine the shape and magnitude of distributed radiological sources.

ACKNOWLEDGEMENTS

REFERENCES

[1] Balraj Singh. Nuclear data sheets for $A = 64$. *Nuclear Data Sheets*, 108(2):197–364, 2007.

[2] Syed M Qaim, T Bisinger, K Hilgers, D Nayak, and Heinz H Coenen. Positron emission intensities in the decay of 64Cu, 76Br and 124I. *Radiochimica acta*, 95(2):67–73, 2007.

[3] Tetsuji Imanaka, Gohei Hayashi, and Satoru Endo. Comparison of the accident process, radioactivity release and ground contamination between Chernobyl and Fukushima-1. *Journal of radiation research*, 56(suppl_1):i56–i61, 2015.

[4] American Elements. Copper. Retrieved Nov. 1, 2021, from https://www.americanelements.com/cu.html.

[5] Richard M Lindstrom, Ronald F Fleming, et al. Neutron self-shielding factors for simple geometries, revisited. *Chem. Anal*, 53:855–859, 2008.

[6] VV Zerkin and B Pritychenko. The experimental nuclear reaction data (EXFOR): Extended computer database and web retrieval system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 888:31–43, 2018.

[7] DW Kneff, LR Greenwood, BM Oliver, RP Skowronski, and EL Callis. Helium production in copper by a thermal three-stage reaction. *Radiation Effects*, 93(1-4):217–220, 1986.

[8] International Tennis Federation. ITF rules of tennis, 2021. Retrieved Sept. 20, 2021, from https://www.itftennis.com/media/4421/2021-rules-of-tennis-english.pdf.

[9] 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.

[10] RT Pavlovsky, JW Cates, M Turqueti, D Hellfeld, V Negut, A Moran, PJ Barton, K Vetter, and BJ 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.

[11] Matthew Newville, Till Stensitzki, Daniel B Allen, Michal Rawlik, Antonino Ingargiola, and Andrew Nelson. LMFIT: Non-linear least-square minimization and curve-fitting for Python. *Astrophysics Source Code Library*, pages ascl–1606, 2016.

[12] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 4th edition, 2010.

[13] John H Hubbell and Stephen M Seltzer. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients 1 keV to 20 MeV for elements $Z = 1$ to 92 and 48 additional substances of dosimetric interest. Technical report, National Inst. of Standards and Technology-PL, Gaithersburg, MD, 1995. Version 1.4.

[14] M.J. Berger, J.H. Hubbell, S.M. Seltzer, J. Chang, J.S. Coursey, R. Sukumar, D.S. Zucker, and K. Olsen. XCOM: Photon cross section database (version 1.5). Retrieved July 5, 2022, from http://physics.nist.gov/xcom.

[15] Weather Underground. Pullman, WA weather history, 2021. Retrieved July 5, 2022, from https://www.wunderground.com/history/daily/us/wa/pullman/KPUW/date/2021-7-9.

[16] Omni Calculator. Air density calculator, 2022. Retrieved July 5, 2022, from https://www.omnicalculator.com/physics/air-density#how-to-calculate-the-air-density.

[17] RS Detwiler, RJ McConn, TF Grimes, SA Upton, and EJ Engel. Compendium of material composition data for radiation transport modeling. Technical report, Pacific Northwest National Lab. (PNNL), Richland, WA (United States), 2021.

[18] L Sissler, R Jones, PG Leaney, and A Harland. Viscoelastic modelling of tennis ball properties. In *IOP conference series: materials science and engineering*, volume 10, page 012114. IOP Publishing, 2010.

[19] Mirion Technologies. Genie™ 2000 Gamma Analysis Software, 2016. Retrieved from https://mirion.s3.amazonaws.com/cms4_mirion/files/pdf/spec-sheets/c47691_genie-2000_gamma_analysis_software_spec_sheet_1.pdf?1557861084 on Feb. 28, 2022.

[20] M Salathe, BJ Quiter, MS Bandstra, JC Curtis, R Meyer, and CH Chow. Determining urban material activities with a vehicle-based multi-sensor system. *Physical Review Research*, 3(2):023070, 2021.

[21] D Hellfeld, MS Bandstra, JR Vavrek, DL Gunter, JC Curtis, M Salathe, R Pavlovsky, V Negut, PJ Barton, JW Cates, et al. Free-moving quantitative gamma-ray imaging. *Scientific reports*, 11(1):1–14, 2021.

[22] THY Joshi, BJ Quiter, J Curtis, MS Bandstra, D Hellfeld, M Salathe, A Moran, J Vavrek, Department of Homeland Security, DOD Defense Threat Reduction Agency, and USDOE. Multi-modal Free-moving Data Fusion (MFDF) v1.0, 2020.

[23] MS Bandstra, D Hellfeld, JR Vavrek, BJ Quiter, K Meehan, PJ Barton, JW Cates, A Moran, V Negut, R Pavlovsky, et al. Improved gamma-ray point source quantification in three dimensions by modeling attenuation in the scene. *IEEE Transactions on Nuclear Science*, 68(11):2637–2646, 2021.

[24] NJ Murtha, LE Sinclair, PRB Saull, A McCann, and AML MacLeod. Tomographic reconstruction of a spatially-extended source from the perimeter of a restricted-access zone using a SCoTSS compton gamma imager. *Journal of Environmental Radioactivity*, 240:106758, 2021.