2D peak fitting for the analysis of radioxenon beta gamma spectra

Jefferson Sesler[1] · James Scoville[1] · Thienbao Carpency[2] · Jordan Armstrong[1] · Robert Hall[1]

Abstract

Atmospheric traces of radioactive xenon, in particular ^{131m}Xe, ^{133}Xe, ^{133m}Xe and ^{135}Xe, can provide "smoking gun" evidence to classify underground nuclear fission reactions. Current software used to quantify isomer concentrations relies on a Region of Interest (ROI) method to sort beta-gamma coincidence counts. This experiences errors when classifying nuclides, especially with metastable nuclides, due to the difficulty of deconvoluting overlapping ROIs and accounting for shifts in detector calibration over time. To address this uncertainty, our technique mathematically models the distinctive peaks in an isomer's beta-gamma spectrum. The function representations are then fitted to measured spectra to determine the concentrations of the primary isomers in the sample. From this proof-of-concept, we hope to create a more precise and accurate system to detect nuclear fission reactions.

Keywords Peak fitting · Least squares fitting · Fermi beta decay · Gaussian distribution · Radioxenon · Beta-gamma analysis · Nuclear treaty monitoring · International Monitoring System

Introduction

The International Monitoring System (IMS), with the Comprehensive Nuclear-Test-Ban Treaty Organization, currently utilizes beta-gamma coincidence counting to detect and quantify four xenon nuclides: ^{131m}Xe, ^{133}Xe, ^{133m}Xe and ^{135}Xe [1–3]. Radioactive isomers of xenon and krypton are both very likely to be generated by thermal neutron fission of uranium and, being noble gases, tend to be difficult to contain compared to particulate or chemically reactive emissions. Xenon isomers are much more useful than krypton ones, however, due to their half-lives of days to several hours. This is long enough to be detectable downwind of a detonation after diffusing through the ground, while also short enough to have low background levels against which new releases are easily noticeable. By examining the concentrations of the four xenon radionuclides, treaty monitoring agencies can use activity ratios to determine if elevated readings came from peaceful nuclear activities, such as power generation or medical isotope production, or if they came from a nuclear weapon explosion [4]. Beta-gamma coincidence counting greatly simplifies concentration measurements by reducing the background noise of beta particles and gamma radiation not associated with xenon radioisotopes [5]. Coincidence counting is accomplished with simple but robust detectors which enclose an atmospheric sample in a sodium iodide shell and a plastic scintillator shell paired with photomultipliers, all enclosed in lead [6]. Simultaneous beta and gamma radiations are recorded, and plotted according to both radiation energies to create a beta-gamma spectrum.

The 2-dimensional (2D) beta-gamma spectra for each of the four radioxenon isomers that we studied are illustrated in Fig. 1. ^{133}Xe and ^{135}Xe both experience beta decay coincident with a gamma ray, while the metastable states ^{131m}Xe and ^{133m}Xe experience a conversion electron decay coincident with a gamma ray [7]. Beta decay releases a beta particle and a neutrino from the nucleus, creating an underdefined three-body system—as a result, beta decays appear on beta-gamma spectra as streaks across many beta energies at a single gamma energy. As shown in Fig. 1, ^{133}Xe has these streaks at a gamma energy of 31.6 keV and at 80 keV, with a maximum beta energy of 346 keV. ^{135}Xe primarily releases 250 keV gamma rays, as well

✉ Thienbao Carpency
thienbao.carpency.ctr@afacademy.af.edu

1 Department of Physics, United States Air Force Academy, USAFA, CO 80840, USA

2 Thienbao Consulting LLC in Coordination with Department of Physics, United States Air Force Academy, USAFA, CO 80840, USA
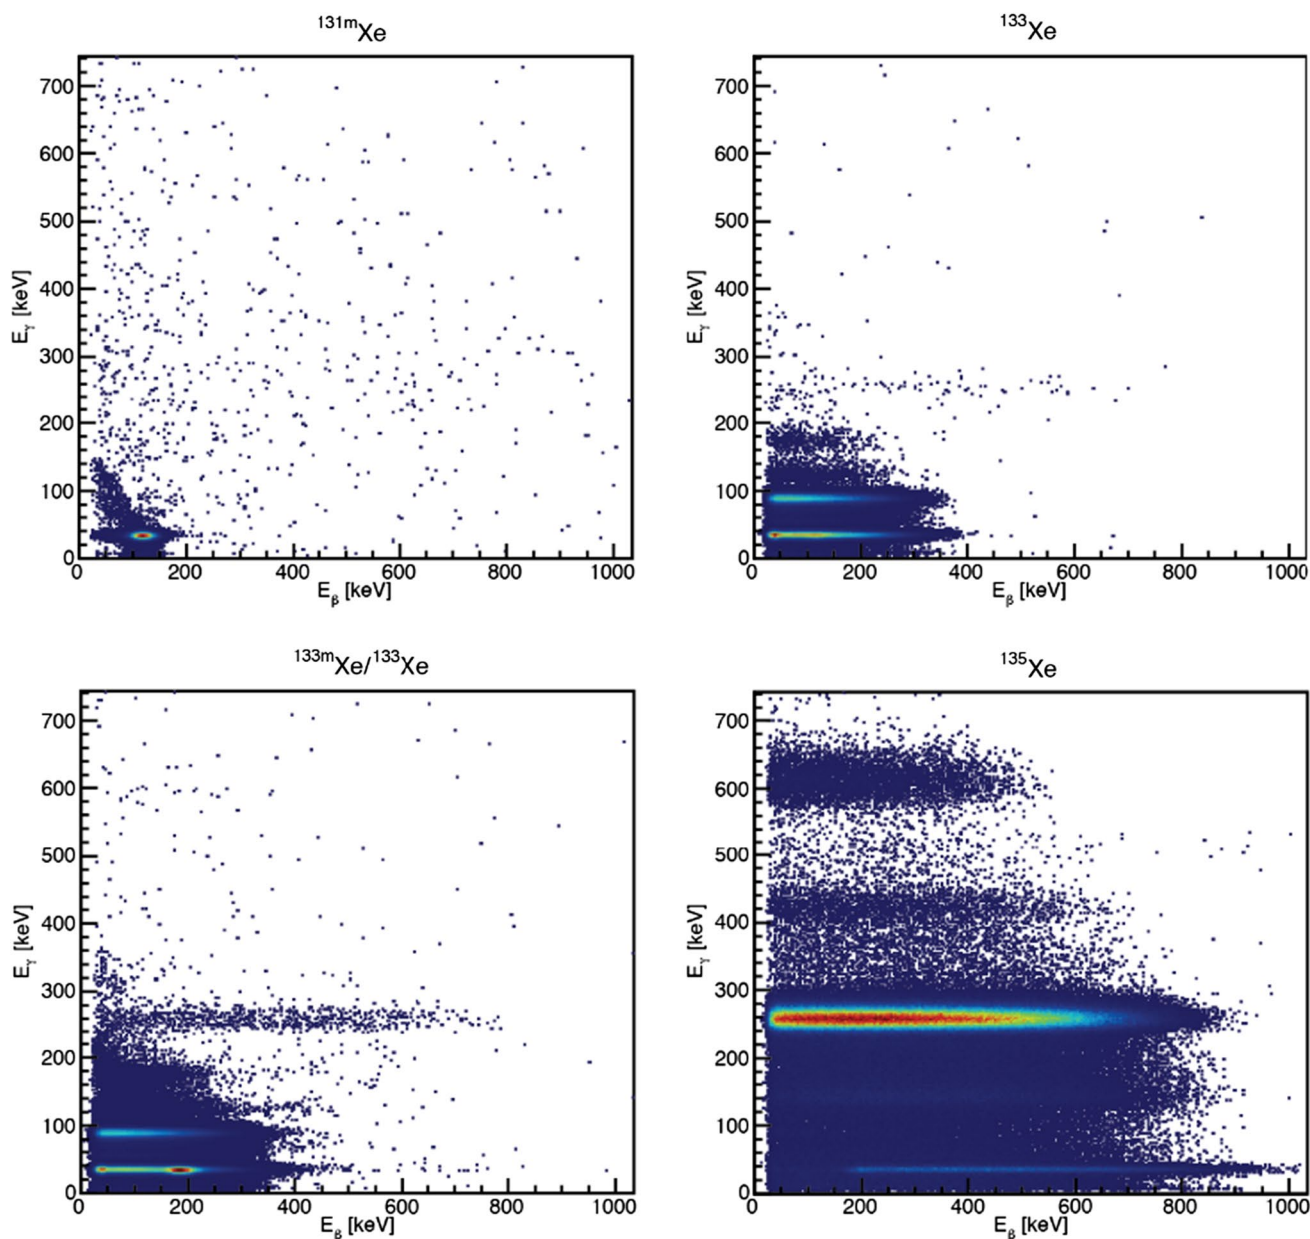
Fig. 1 The 2D beta-gamma energy spectra for ^{131m}Xe, ^{133}Xe, ^{133m}Xe & ^{133}Xe, and ^{135}Xe [8]

as a few 31.6 keV rays, with a maximum beta energy of 910 keV [7]. Internal conversion decay does not release a neutrino, and so both the released electrons and gamma rays have a definite energy: 130 keV and 30.4 keV for ^{131m}Xe, and 199 keV and 30.4 keV for ^{133m}Xe. Note that ^{133m}Xe cannot be isolated from its decay product ^{133}Xe, however its unique spectrum is identifiable by comparing the pure ^{133}Xe spectra to the ^{133m}Xe/^{133}Xe combined spectra in Fig. 1.

The current software used to classify ^{131m}Xe, ^{133}Xe, ^{133m}Xe and ^{135}Xe relies on Region of Interest (ROI) counting. This subdivides the beta-gamma spectrum into several

regions roughly corresponding to the activity from each nuclide. The number of counts in each region are summed, and counts in regions encompassing multiple interfering isomers are deconvoluted, to calculate the concentration of each isomer. While relatively simple, this method can be imprecise due to overlapping regions of isomer activity, where deconvolution becomes especially sensitive to gain shifts in the energy calibration of the detector over time. The meta-stable isomers' spectra are particularly vulnerable to these deconvolution errors, since their activity is completely overlapped by other nuclides' spectra [5, 9].

Several groups have developed various analysis techniques for radio-xenon beta-gamma spectra. The Nuclear Explosion Monitoring and Policy program at Pacific Northwest National Laboratory (PNNL) improved upon the ROI method by changing the number of regions of interest. By changing the number of regions, they are able to more easily deconvolute overlapping regions, thus giving a more accurate result in comparison to the normal ROI method, but this method remains vulnerable to gain shifts in energy calibration [5]. Researchers at The University of Texas at Austin developed another improvement to the ROI method known as the Spectral Deconvolution Analysis Tool (SDAT). SDAT utilizes software that deconvolves the 3-D sample spectra into the most probable combination of pure radio-isomer signals using the non-negative least-squares method. This method results in a better ability to resolve inferences and improves counting statistics, but, similar to PNNL's improvement, requires complex corrections to correct for gain shifts [10–13].

In this work, we aim to address this problem by modeling beta-gamma spectra using mathematical functions, which can be easily and automatically adjusted to account for gain shifts. The energy distributions of two-dimensional spectral histograms can be modeled by two-dimensional functions, allowing the distributions to be summarized by function parameters. Constrained least-squares fitting can then fit multiple functions onto the peaks from sample data, adjusting the location and scaling parameters of the function to match detector gain. Once fitted, function amplitude replaces count number as the proxy for nuclide activity, and with sufficient calibration can be used to compute nuclide concentrations in the sample.

Theory

The two-dimensional energy distributions for each spectral feature is the product of a gamma probability density function (PDF) ($n_\gamma(E_\gamma)$) and an electron PDF ($n_e(T)$);

$$n_{2D\,feature}(T, E_\gamma) = A\, n_e(T)\, n_\gamma(E_\gamma) \tag{1}$$

where A is the feature amplitude, E_γ is the gamma ray energy, and T is the electron kinetic energy. The functional form of the electron PDF varies depending on the decay type, whether conversion electron or beta decay. The one-dimensional PDFs are normalized so that the area under each distribution is one. One or more 2D spectral feature PDFs sum to fully model a particular isomer

$$n_{isomer}(T, E_\gamma) = \sum_{2D\,features} n_{2D\,feature}(T, E_\gamma) \tag{2}$$

and the four isomer PDFs sum to describe the full xenon spectra.

Each 2D spectral feature PDF contains several fitting parameters, determining the peak location, width, and amplitude, which are allowed to vary when calibrating the detector response to a particular isotope. After initial calibration, we improve the algorithm's performance by reducing the number of fit parameters to only those needed to adjust for detector gain shifts and changes in isomer concentrations. Isomer amplitudes are used to determine nuclide concentrations since they are proportional to the decay counts, and thus the concentration, of each nuclide. Before describing how gain shifts are handled, first we detail the functional forms of the 1D PDFs.

We used three 1D PDFs, with shapes based on the decay mechanisms underlying the data they represent. The first of these functions is the Gaussian Distribution, Eq. (3), used to model gamma ray energy distributions for all four isomers as well as the energies of conversion electrons produced by ^{131m}Xe and ^{133m}Xe. Each isomer produces gamma rays and/or conversion electrons within a narrow energy range, which is then broadened by the detector resolution. In accordance with the Central Limit Theorem, [14] these random errors are distributed normally around the true value, with an energy distribution $n(E)$ given by the expression

$$n(E) \propto e^{-\frac{1}{2}\left(\frac{E - E_{peak}}{\sigma}\right)^2} \tag{3}$$

in terms of the peak energy E_{peak} and the standard deviation, or width, σ. Figure 2 exhibits two such Gaussian functions (orange lines) mapped closely to the conversion electron and gamma ray spectra of ^{131m}Xe (blue dots).

The second and third energy distributions we use are described by Fermi beta decay theory and model the variable energies of electrons and neutrinos in the beta decays of ^{133}Xe and ^{135}Xe. The theory predicts the statistical distribution of decay products' momenta using Fermi's Golden Rule, which says the decay probability is proportional to daughter particles' density of states

$$dn = dn_e dn_\nu \propto p_\beta^2 dp_\beta\, p_\nu^2 dp_\nu \tag{4}$$

where dn are the number of states in a momentum interval dp and p_β and p_ν are the beta particle and neutrino momenta respectively [15]. From the Einstein energy-momentum relation[1] $E^2 - p^2 = m^2$ we have $p\,dp = E\,dE$ and we can re-write Eq. (4) as

$$dn \propto (E_\beta p_\beta E_\nu p_\nu)\, dE_\beta\, dE_\nu . \tag{5}$$

---

[1] We use natural units wherein energy, mass, and momentum all have units of energy.

Fig. 2 The beta-gamma spectrum 2D histogram of a ^{131m}Xe conversion electron emission, with the 2D fit overlaid with contour lines. The 1D electron and gamma Gaussian fits (orange lines) overlay the total beta and gamma counts (blue dots) on the axes, with the residual in the electron fit also plotted. (Color figure online)
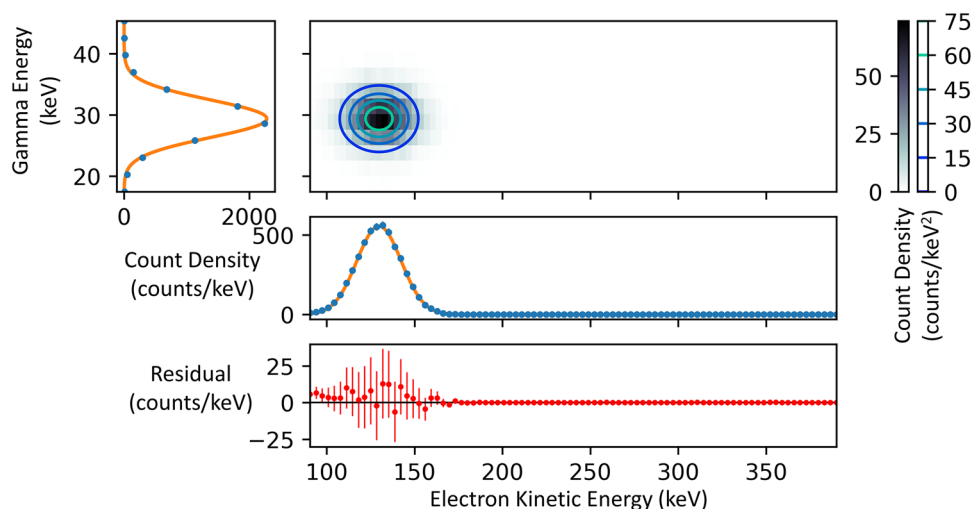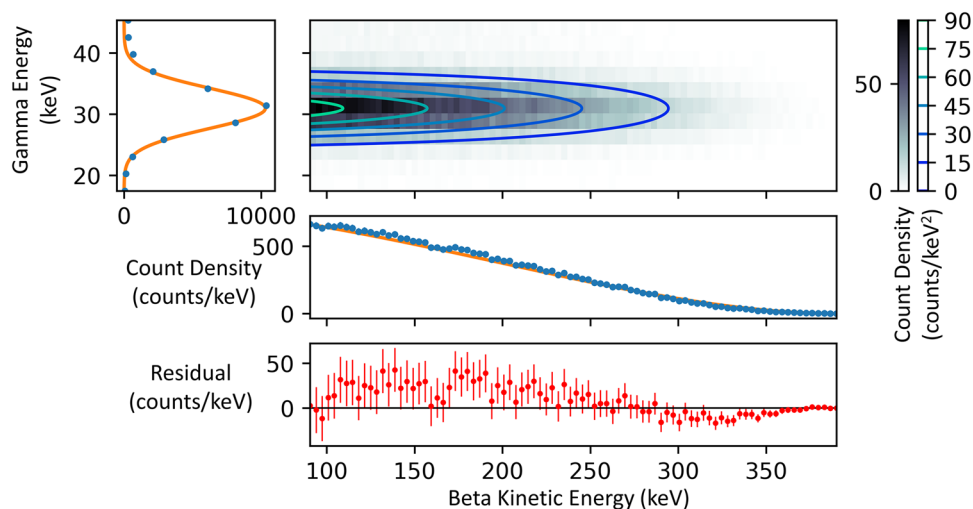
Fig. 3 The beta-gamma spectrum 2D histogram for the $E_\gamma = 30$ keV spectral feature in ^{133}Xe beta decay, with the 2D fit overlaid with contour lines. The 1D Fermi beta fit and gamma Gaussian fit (orange lines) overlay the total beta and gamma counts (blue dots) on the axes, with the residual in the beta fit also plotted. (Color figure online)

The term proceeding the energy differentials is called the *statistical factor*. It can be expressed entirely in term of the beta particle's measured kinetic energy T using Einstein's energy-momentum relation along with the expressions $E_\beta = T + m_\beta$ and $Q = T + E_\nu$, where Q is the total decay kinetic energy.

Two additional terms modify the simple distribution given in Eq. (5). The first is the *shape function*, $S(T)$, which includes the nuclear matrix element. For beta decays that are well described by the "allowed approximation" the shape function is constant. The second term, called the *Fermi function*, describes how the shape of the beta kinetic energy distribution is modified by Coulomb interactions with the daughter nucleus. Taken together, the beta particle PDF, $n(T)$, is given by

$$n_\beta(T) \propto E_\beta p_\beta E_\nu p_\nu F(T)S(T) \tag{6}$$

where the energy and momenta are implicit functions of T [15].

The beta decay of ^{133}Xe is an allowed transition ($S(T) = 1$) and the Fermi function is well approximated by $F(T) = E_\beta/p_\beta$. The negligible mass of the neutrino means its energy and momentum are nearly equal, thus we may simplify Eq. (6) to

$$n_\beta(T) \propto E_\beta^2 E_\nu^2, \tag{7}$$

or in terms of T as

$$n_\beta(T) \propto (T + m_\beta)^2 (Q - T)^2. \tag{8}$$

This curve is fitted to the beta decay of ^{133}Xe in Fig. 3.

Finally, we have the beta decay of ^{135}Xe, which experiences an allowed transition at a much higher cutoff energy Q than ^{133}Xe. The higher-energy beta particles will sometimes "punch-through" through the plastic scintillator detectors, depositing only a portion of their energy, which results

Fig. 4 The beta-gamma spectrum 2D histogram of ^{135}Xe, with the 2D fit overlaid with contour lines. The 1D Fermi beta fit and gamma Gaussian fit (orange lines) overlay the total beta and gamma counts (blue dots) on the axes, with the residual in the beta fit also plotted. (Color figure online)
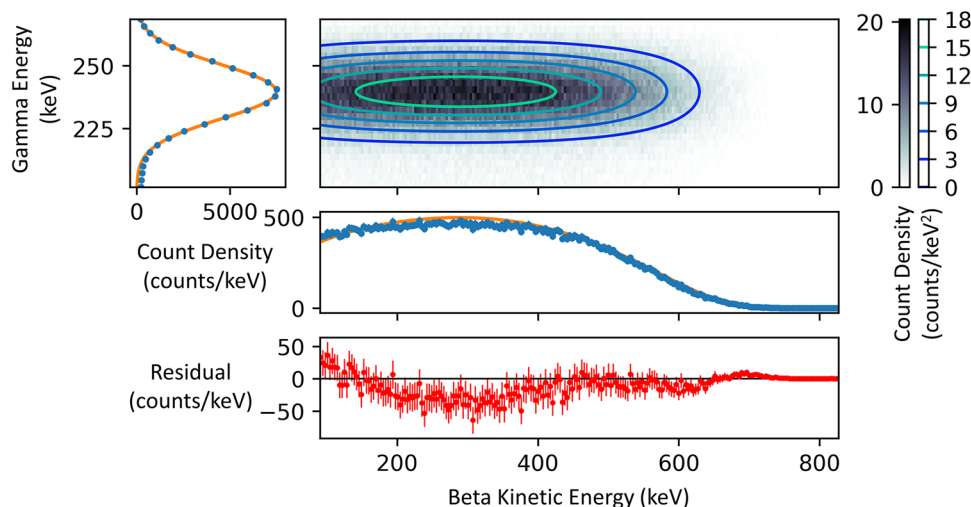
in detector undercounts of high-energy particles and over-counts of low-energy particles. Rather than trying to correct for this altered shape, we noted that the resulting "S-curve" was close to the distribution predicted by Fermi theory for first forbidden transitions. From Eq. (6), we used a Fermi function $F(T) = 1$ and the first-forbidden shape function $S(T) = p_\beta^2 + p_\nu^2$ to generate a ^{135}Xe beta energy distribution model of

$$n_\beta(T) \propto E_\beta p_\beta E_\nu^2 \left(p_\beta^2 + p_\nu^2 \right). \tag{9}$$

Substituting the definitions for total beta energy E_β and momentum p_β, and neutrino energy E_ν and momentum p_ν, in terms of beta mass m_β, beta kinetic energy T, and total decay kinetic energy Q gives

$$n_\beta(T) \propto \left(T + m_\beta \right) \sqrt{T^2 + 2Tm_\beta} (Q - T)^2 \left(T^2 + 2Tm_\beta + (Q - T)^2 \right). \tag{10}$$

While not an exact fit, Fig. 4 shows this curve reliably following the measured beta energy distribution of ^{135}Xe.

We substitute these 1D PDFs into Eq. (1) to arrive at the 2D spectral feature PDFs. In particular, we arrive at

$$n_{CE}(T, E_\gamma) = A \cdot \left(N_{CE} e^{-\frac{1}{2} \left(\frac{T - T_{CE\,peak}}{\sigma_{CE}} \right)^2} \right)$$
$$\cdot \left(N_\gamma e^{-\frac{1}{2} \left(\frac{E_\gamma - E_{\gamma\,peak}}{\sigma_\gamma} \right)^2} \right), \tag{11}$$

$$n_\beta(T, E_\gamma) = A \cdot \left(N_\beta (T + m_\beta)^2 (Q - E_\beta)^2 \right)$$
$$\cdot \left(N_\gamma e^{-\frac{1}{2} \left(\frac{E_\gamma - E_{\gamma\,peak}}{\sigma_\gamma} \right)^2} \right), \tag{12}$$

and

$$n_\beta(T, E_\gamma) = A \cdot \left(N_\beta (T + m_\beta) \sqrt{T^2 + 2Tm_\beta} (Q - T)^2 \right.$$
$$\cdot \left(T^2 + 2Tm_\beta + (Q - T)^2 \right) \Big)$$
$$\cdot \left(N_\gamma e^{-\frac{1}{2} \left(\frac{E_\gamma - E_{\gamma\,peak}}{\sigma_\gamma} \right)^2} \right), \tag{13}$$

where the N's in Eqs. (11), (12), and (13) are normalization coefficients chosen so that the area under each 1D distribution is one.[2] This normalization ensures that the function amplitudes (A) are approximately equal to the number of counts in the histogram bins they cover. The conversion electron 2D peak function, Eq. (11) thus has five parameters for fitting: the peak electron and gamma energies ($T_{CE\,peak}$ and $E_{\gamma\,peak}$) giving the "location" of the peak on the beta-gamma spectrum, the electron and gamma peak widths (σ_{CE} and σ_γ), and the amplitude (A) giving the "height." Likewise, the beta decay 2D PDFs, Eqs. (12) and (13), have four fitting parameters: the two electron parameters are replaced by the maximum beta kinetic energy (Q), which defines the maximum extent of the Fermi distribution, and serves as the "location" parameter. ^{133}Xe and ^{135}Xe generate gamma rays at multiple energies creating multiple distinct spectral features: we combine these into a single isomer model by summing the features (with amplitudes in fixed ratios, depending on the particular isomer's spectra) as in Eq. (2) (Fig. 5).

Our algorithm fits the isomer PDFs to the spectral data via least squares fitting. When an isomer is absent from a sample spectrum, the algorithm should fit that isomer's

---

[2] For Gaussian functions the normalization coefficient is the usual $N = \frac{1}{\sigma \sqrt{2\pi}}$. For the beta distributions the functional form is more complicated but can be determined by integration.
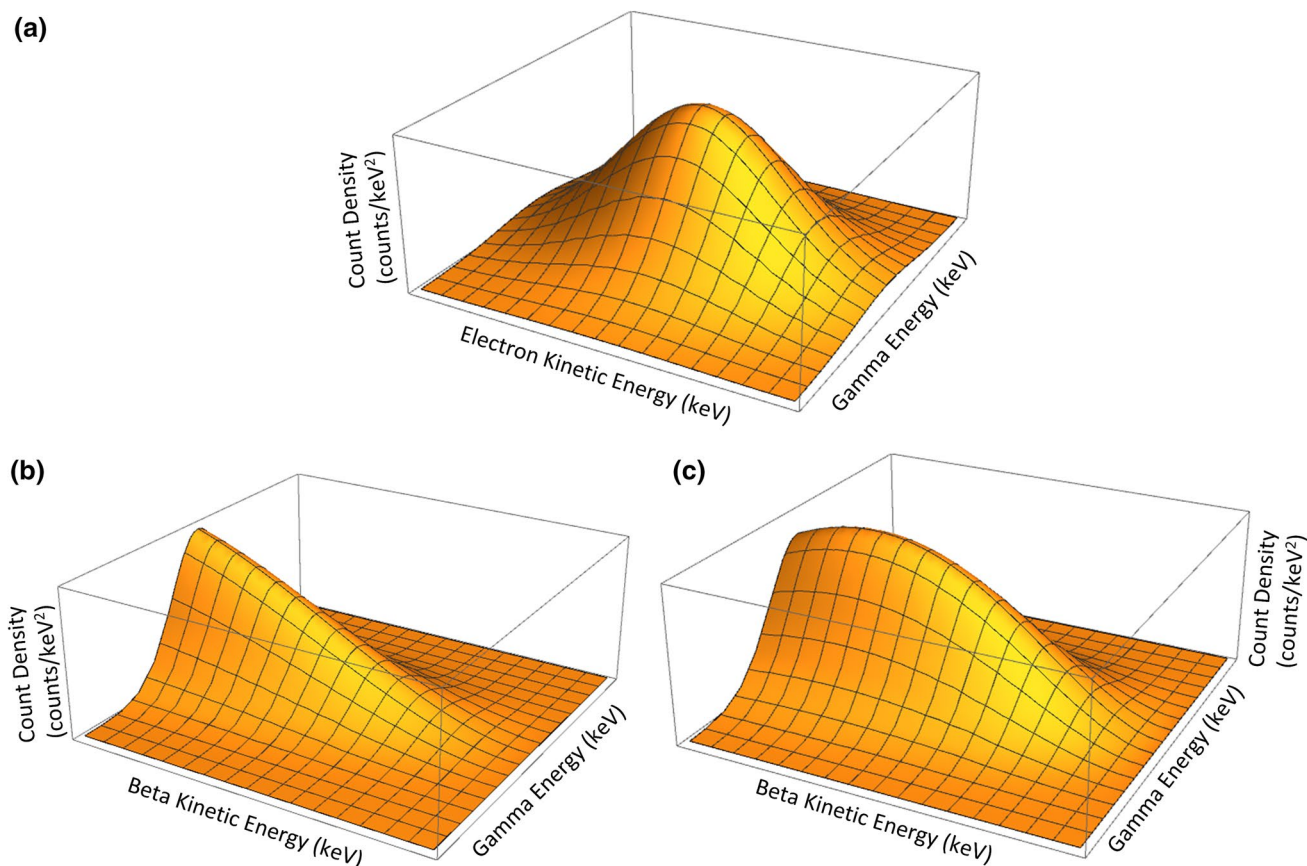
(a)

(b)

(c)

Fig. 5 Graphics visualizing the shapes of the 2D PDFs for **a** a conversion electron peak, **b** a ^{133}Xe beta decay peak, and **c** a ^{135}Xe beta decay peak

corresponding PDF to the data by pushing its amplitude to zero, keeping the other function parameters unchanged. While the function may not fit with exactly zero amplitude, due to uncertainty in the measurement, we look for the amplitude to be small enough that it does not induce significant error in the nonzero PDF models. This requires adding constraints to the fitting routine, as without them, the peak functions would "wander." This wandering could fit peaks to background noise, giving nonzero amplitudes; or it could fit to portions of spectral features from other nuclides, changing both amplitudes; or the functions could move outside the spectrum where it would behave unpredictably. To address this, we constrained the function parameters by linking the location and width parameters for all four isomer PDFs into four fit parameters describing scale and translation in both dimensions (electron energy and gamma energy) for the entire spectrum, as shown in Fig. 6. The algorithm takes initial values for each 2D PDF, which were generated by an unconstrained least-squares fit of each PDF to only its corresponding feature. This initialization calibration must be done for each detector, due to each detector's unique biases. Then during subsequent full-spectrum fitting, the model's electron-dimension

parameters can be scaled, but they must all be scaled by the same factor. Likewise, the algorithm can shift all the electron-dimension location parameters by the same amount, and it can adjust the gamma-dimension parameters in the same manner. Thus the algorithm can move or squeeze the entire function mapping to best fit the data, easily adjusting for detector gain shifts, but no single peak's location or width can change relative to the others.

We incorporated other ancillary constraints to improve performance. We prohibited fitting parameters for PDFs from becoming negative where necessary. We also incorporated a very slight cost to changing the four constrained fit parameters, by treating the difference between each parameter and its initial value as an error to be minimized alongside the difference between each data point and its modeled value. This was weighted enough to prevent wandering from the initial values when fitting to a completely empty spectrum, while also allowing unencumbered adjustments to fit to shifted data. We then used a standard Python least-squares error optimizing function to find the best values of the four scaling/translation parameters and the four isomer amplitudes that fully define the 2D peak functions. Once fitted to the data, these isomer function amplitudes
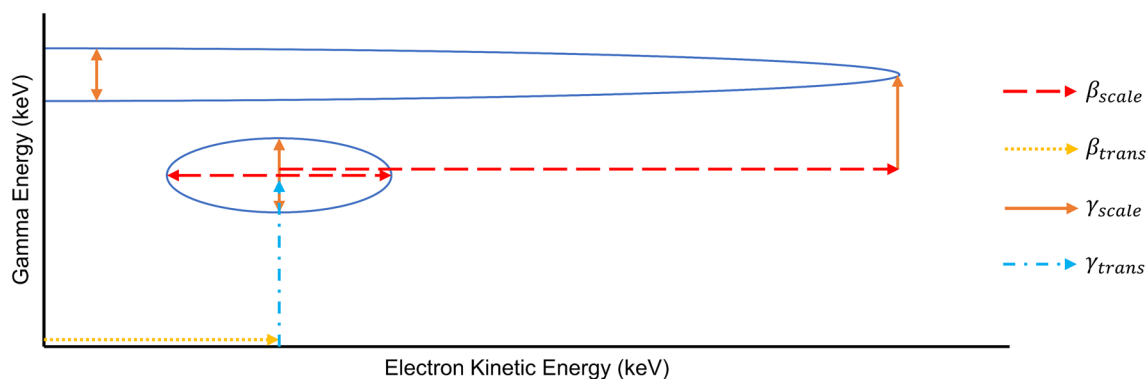
Fig. 6 An illustration of the concept behind constraining all isomer PDFs to four fit parameters. The scaling parameters control the size of and distance between peaks, while the translation parameters shift all the peaks together

are proportional to the concentrations of the nuclides in the detector measurements.

Results

We evaluated the effectiveness of the peak-fitting routine and the constraints applied to it with two tests. The first was designed around the limited spectra from available from real detectors, while the second used a set of simulated data to expose the algorithm to a wide variety of scenarios, allowing a much more rigorous baseline than would be practical from experimental data alone.

The first test assessed the self-consistency of fitting in three cases: in the first (control) case we fit data from a single isomer decay spectra using only the corresponding isomer PDF, in the second (experiment) case we fit all four isomer PDFs to the same single-isomer data, before comparing the results, and in the third (experiment) case we repeated the experiment with a four-isomer mixture. We used spectral data obtained from the 21st Reconnaissance Squadron at the Air Force Technical Applications Center, which was generated by four detectors, with one sample of isolated ^{131m}Xe, ^{133}Xe, and ^{135}Xe per detector. Since we did not have detector samples of ^{133m}Xe, we complemented the data with a simulated spectra, allowing us to study the interactions of all four nuclides. In the control case (single isomer data fitted with a matching single isomer model), the resulting PDF fit amplitudes were taken as the baseline against which the experiment cases (multiple isomer models fitted to data) were compared. Our objective with the first experiment, fitting all four isomer PDFs to data from a single isomer, was to determine if the models interfered with each other during fitting. In proper operation, the peak corresponding to the nuclide present would be identical to the control and the amplitudes of the other three models would be close to zero. Any error in the primary isomer amplitude could then be taken as an indication of interference from another model. For the second experiment, we created four spectra (one for each detector) with a synthetic mixture of all four isomer spectra by summing together the pure isomer data for every detector, and fit the combined isomer PDFs to them. Since the counts from each isomer were the same as the separated controls, each isomer amplitude should be the same as well; our goal was to see if the isomers' overlapping spectral features would be properly deconvoluted to give the same amplitudes. All three cases from this test are illustrated by Fig. 7 alongside the raw data.

The results from the single-isomer tests are given in Table 1, those from the mixed-isomer tests are in Table 2. Across all four detector data sets, the (control) single isomer data fits to a single matching PDF were almost indistinguishable from the (experiment) multi-model fits to the same data, demonstrating success in using the least-squares error fit to isolate the proper isomer PDF and push the other three to zero amplitude. In this way, the algorithm is simultaneously identifying what isomers are present and which are absent, and measuring the concentration of those present. Measuring all four nuclides at once in the synthetic mixture introduced slightly higher amplitude errors, although most measurements were still well within 1% accuracy. While clearly successful, it should be noted that due to having access to only one sample of each isomer from each detector, the baseline peak location and width calibration for each detector had to be done using the same samples the program was tested on. We addressed this in the simulated-data experiment by obtaining all samples using the same simulator settings, requiring only one calibration.

For our second set of experiments, we used the BGSim spectra simulator provided by Pacific Northwest National Laboratory to generate 625 beta-gamma spectra. [4] These encompassed all permutations of the four isomers at 5 different concentrations: 0, 10,000, 1,000,000, 3,000,000, and 10,000,000 atoms in a 1.3 cubic centimeter sample. For
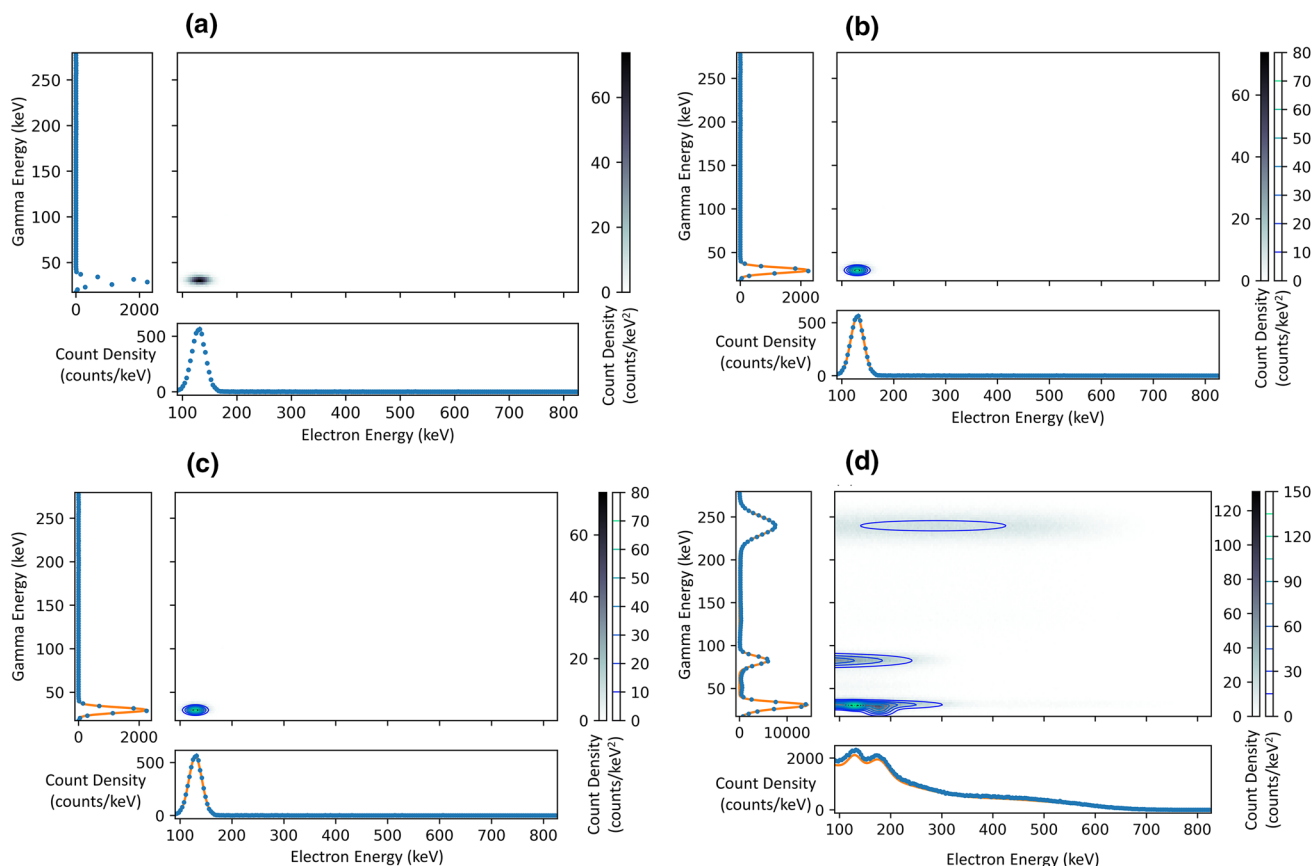
Fig. 7 Four plots demonstrate stages of the self-consistency experiments. The plot (**a**) shows a beta-gamma spectrum of pure ^{131m}Xe from a detector, with the gamma and the electron distributions below. Plot (**b**) overlays the ^{131m}Xe PDF fitted to the data, the control case. On the 2D spectra, the fitted model is shown with contour lines, with color corresponding to height; on the 1D plots, the model is shown by the orange lines as before. Plot (**c**) shows all four isomer PDFs fitted to the data: note that it is nearly identical to the control. Plot (**d**) shows the synthetic mixture with all four models fitted to it

Table 1 Amplitude errors in primary isomer of a four-PDF fit to single-isomer data, using single-isomer PDF fits to the same data as a baseline

Isomer	Det 0 (%)	Det 1 (%)	Det 2 (%)	Det 3 (%)
^{131m}Xe	0.191	0.362	0.267	0.135
^{133}Xe	0.125	0.269	0.309	0.0021
^{135}Xe	0.000079	0.00031	0.000050	0.0022

The low errors indicate the algorithm's ability to successfully identify and quantify pure isomer samples while rejecting isomers not present

Table 2 Amplitude errors in the fits of four isomer PDFs to spectral data from a mixture of all four isomers; single-isomer PDF fits to single-isomer data served as a baseline

Isomer	Det 0 (%)	Det 1 (%)	Det 2 (%)	Det 3 (%)
^{131m}Xe	1.05	0.185	2.76	3.22
^{133m}Xe	0.253	0.0571	1.28	1.59
^{133}Xe	0.147	0.329	0.0533	1.02
^{135}Xe	0.285	0.322	0.133	0.510

The four-isomer mixture data sets were generated by summing the pure isotope spectra prior to analysis. We did this separately for each detector, using simulated data for ^{133m}Xe. The low errors indicate the algorithm's ability to successfully deconvolute the spectra of all four isomers in each of the detectors tested

any particular isomer concentration there were $5^3 = 125$ spectra from permutations of the other three isotopes at each of the five concentrations. We used four pure-isomer data sets (one for each isomer) with 10,000,000 atoms each to calibrate the algorithm to the simulated detector response, both in terms of PDF shapes and amplitudes. We then analyzed all 625 spectra using a four-isomer fit, and

calculated the percent error in each isomer concentration prediction from the value used to generate the simulated data.

Results from the simulated-data experiment are summarized in Tables 3 and 4. There are 500 different percent

Table 3 Simulated dataset fitting results, arranged by true number of atoms

Isomer	True count	Calculated count $(\times 10^3)$			Percent error (%)						
		Median	Mean	SD	Outlier min	Min	Median	Max	Outlier max	Mean	SD
^{131m}Xe	0	0	6.09	17.9							
	10K	0	9.24	18.4		0.441	100	928		**128**	133
	1M	967	938	85.8	0.0022	0.0120	3.68	30.6		7.10	7.85
	3M	2930	2910	86.9	0.0075	0.0917	2.19	9.96		3.22	2.80
	10M	9880	9830	112		0.0637	1.19	4.52		1.66	1.12
^{133m}Xe	0	7.57	37.2	53.8							
	10K	19.7	44.7	55.5	0.0242	0.235	100	1590		395	522
	1M	997	1020	60.4		0.0811	2.33	16.8		4.39	4.74
	3M	2970	3010	74.6	0.0042	0.0587	1.33	7.01		1.88	1.64
	10M	9890	9930	86.9		0.0126	1.15	1.81		1.08	0.386
^{133}Xe	0	13.6	30.0	40.6							
	10K	22.5	35.2	36.7	0.0413	1.38	125	1630		305	325
	1M	982	971	100.		0.0178	3.55	22.7	86.7	5.54	8.82
	3M	2910	2900	75.1		0.0105	3.20	9.91		3.44	2.29
	10M	9640	9640	107		1.08	3.59	6.23		3.57	1.07
^{135}Xe	0	0	0.222	0.941							
	10K	8.63	6.74	4030	0.0022	0.0579	14.7	100		35.4	37.8
	1M	994	984	79.8		0.0201	0.576	3.89	89.9	1.57	7.97
	3M	2990	2960	250	0.0020	0.0229	0.363	1.96	66.8	1.50	8.33
	10M	9970	9920	623	0.0001	0.0079	0.278	0.720	70.2	0.836	6.23

Each row summarizes the 125 spectra with the indicated isomer at the indicated concentration, across all other permutations of concentrations of the other three isomers. Thus, each of the 625 spectra are represented once in every isomer grouping. Each row indicates the minimum, median, maximum, mean, and standard deviation of the isomer concentration percent errors calculated in all 125 spectra. Outliers are separated from the maximum and minimum columns, to better summarize the data as a whole

errors associated with each isomer arising from the 125 spectra for each of the four nonzero isomer concentrations. Table 3 summarizes the concentration accuracy for each isomer grouped by the isomer concentration. It lists the minimum, median, maximum, mean, and standard deviation of the 125 isomer concentration percent errors. As can be expected, we see that higher concentrations can be detected with greater accuracy. This is partially due to the better-defined spectral features of high-concentration isomers. More often, however, the increased error which sometimes occurs at lower concentrations is associated with mixtures wherein one isomer's concentration is less than another isomer's concentration by an order of magnitude or more, drowning out the smaller isomer's signal.

To further investigate cases with highly unbalanced isomer concentrations, we summarized the same results re-grouped by isomer concentration ratio in Table 4. For this, each beta-gamma spectrum was classified for each isomer according the smallest ratio between the concentrations of that isomer and the other three. For example, all spectra in the "^{131m}Xe $\geq 2 : 1$" grouping have ^{131m}Xe at a greater concentration than any other isomer, by at least a 2:1 ratio, whereas spectra in the "^{131m}Xe $\geq 1 : 1000$" grouping each have at least one isomer at a concentration over 10 times, but less than or equal to 1000 times, that of ^{131m}Xe. In this formulation, we see that most of the high-error cases were at very small concentration ratios, where the isomer being measured is outnumbered by others by at least a hundred to one. Outside of these cases, the least-squares fit reliably predicts concentrations within about 5% of the true value, frequently achieving pinpoint accuracy below 1% but occasionally missing by up to 20–30%.

There are three special cases of increased error apparent in the simulated dataset results worth discussing. The first is the presence of significant outliers, single fits with an amplitude error much smaller or much larger than any other in their group. These have been indicated in both data tables by showing them alongside the second-most extreme in their respective groups. The root cause of these outliers is being investigated, but they occur only in very specific situations. The second special case is the low activity of ^{131m}Xe compared to the other three nuclides, which causes it to produce significantly fewer decay counts, and thus smaller spectral features, than the other nuclides at the same concentration. This is the primary cause for the higher ^{131m}Xe errors when evaluating samples by their concentration. Finally, ^{133}Xe

Table 4 Simulated dataset fitting results, arranged by isomer concentration ratios

Isomer	Ratio	Samples	Percent error (%)						
			Outlier min	Min	Median	Max	Outlier max	Mean	SD
^{131m}Xe	≥ 2:1	100	0.0075	0.0637	0.991	3.41	17.6	1.25	1.76
	≥ 1:1	124		0.231	2.73	27.3		2.84	3.07
	≥ 1:10	159	0.0022	0.0120	4.90	30.6		6.99	7.09
	≥ 1:1000	117		4.04	100	928		136	134
^{133m}Xe	≥ 2:1	100		0.0399	1.20	1.95	9.07	1.15	0.886
	≥ 1:1	124		0.0126	1.17	2.66	9.04	1.20	0.874
	≥ 1:10	159	0.0041	0.0811	2.59	16.8		4.29	4.29
	≥ 1:1000	117		0.396	123	1590		422	529
^{133}Xe	≥ 2:1	100		2.17	3.67	6.41	16.0	3.89	1.46
	≥ 1:1	124	0.0413	0.544	3.37	15.7		3.58	1.91
	≥ 1:10	159		0.0105	2.98	22.7	86.7	4.94	8.08
	≥ 1:1000	117		1.38	149	1630		325	326
^{135}Xe	≥ 2:1	100	0.00014	0.0147	0.242	2.87	89.9	1.83	11.1
	≥ 1:1	124		0.0079	0.348	15.1	70.2	1.19	6.42
	≥ 1:10	159		0.0305	0.616	3.89	66.7	1.25	5.26
	≥ 1:1000	117	0.0022	0.0579	18.1	100		37.5	38.2

Each row summarizes all spectra where the indicated isomer is within the specified range of ratios with respect to the other three isomers. For example, all spectra in the "^{131m}Xe ≥ 2:1" grouping have ^{131m}Xe at a greater concentration than any other isomer, by at least a 2:1 ratio, whereas spectra in the "^{131m}Xe ≥ 1:1000" grouping each have at least one isomer at a concentration over 10 times, but less than or equal to 1000 times, that of ^{131m}Xe. The number of spectra that fall into each ratio group is indicated. Outliers are separated from the maximum and minimum columns, to better summarize the data as a whole

experiences much higher error in fits to simulated data because the BGSim program produces ^{133}Xe spectra shaped differently from experimental results [4]. Figure 8 illustrates this problem, showing the shape of the beta particle energy distribution from experimental data and from the BGSim simulator. This causes difficulty fitting the Fermi distribution curve to the differently-shaped simulator distribution, and also induces error in ^{131m}Xe and ^{133m}Xe fits as the algorithm uses their models to reduce the ^{133}Xe count discrepancy. Although this problem could be addressed by changing the ^{133}Xe shape, we chose to not pursue this since the problem is not present in actual experimental data, where the Fermi PDF matches the data quite accurately. In future work, we hope to obtain a far more extensive set of experimental data to test the algorithm against, hopefully eliminating this issue.

Conclusion and Future Work

Our work aimed to accurately measure radioxenon nuclide concentrations using a constrained least-squares error fitting of two dimensional probability density functions to the beta-gamma spectra generated by detectors. The algorithm was built to identify and robustly handle gain shifts in the detector, and to accurately deconvolute interfering signals from each of the isomers ^{131m}Xe, ^{133m}Xe, ^{133}Xe, and ^{135}Xe. Against a limited set of real detector results, the algorithm performed

remarkably well, reliably predicting the correct isomer function amplitude within 1-4% and usually within 0.5%. The program also made accurate predictions for a much larger and more challenging set of simulated beta-gamma-spectra, with errors in most scenarios of interest below 5%. These results indicate that with further development to address error trends and outlier scenarios, and to incorporate more mathematical tools for real-life spectrum interpretation, the peak fitting methods proposed here can be a viable basis for future radioxenon measurement efforts.

In the immediate future we will be continuing to evaluate the results from the large-batch simulated data experiments, to understand the specific behavior of the least-squares fit in the high-error outlier cases, to identify trends correlating with higher amplitude error, and to develop corrective measures in the functions and constraints. We are also working with Pacific Northwest National Labs and the Air Force Technical Applications Center to acquire more detector test data against which to run our proof-of-concept program, as well as samples with low radioxenon counts, high background noise, or high detector gain to challenge the versatility of the program. In this way, we hope to continue incremental improvements of the algorithm and to drive down its error statistics.

Improvements to the PDF fitting algorithm are also needed. The current algorithm minimizes the sum of the squares of the difference between the PDF and the data.

Fig. 8 These plots illustrate the incorrect shape of the simulated ^{133}Xe beta energy distribution. Plot (**a**) is the beta-gamma spectrum from a detector test: note the nearly linear beta distribution in the center axes, which is well matched by the fit function in orange. Plot (**b**) shows the simulated spectrum, fitted only with the ^{133}Xe model. Note the distinctive curve in the beta distribution, which is not followed by the more linear fit function. Finally, plot (**c**) shows the same data fitted with all isomer functions, demonstrating how the algorithm attempts to "fill in" the gap between the model and the data by incorrectly adding the ^{133m}Xe function, creating a lump in the orange fitted distribution. This behavior is absent when fitting to actual experimental spectra due to the better beta function fit
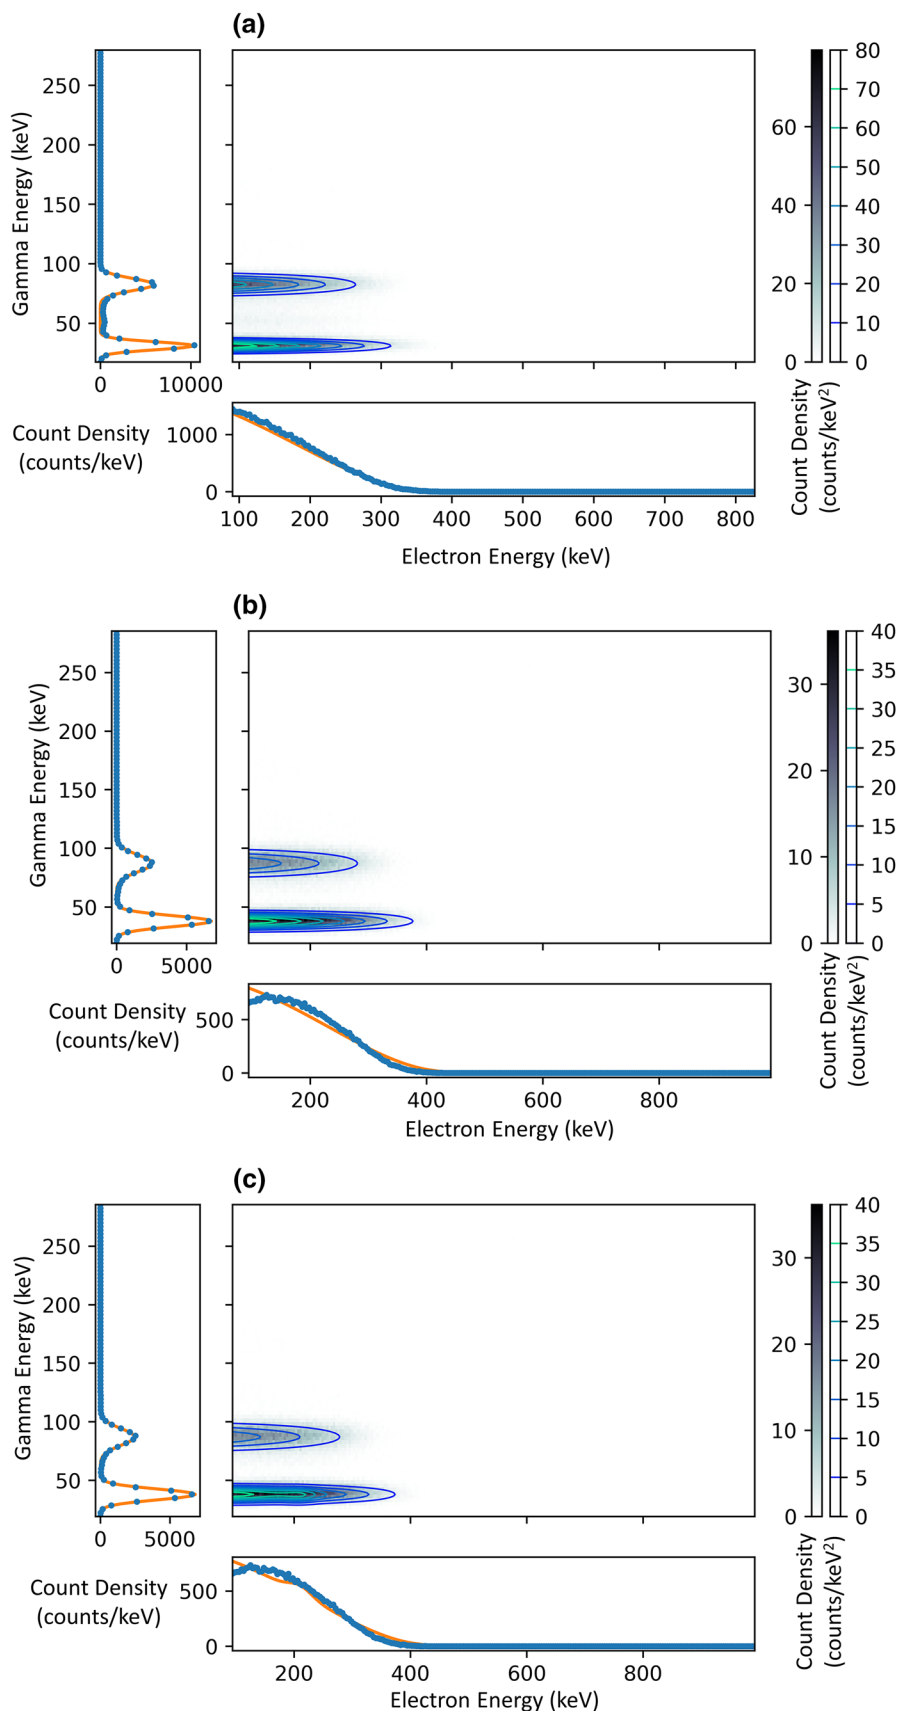
Data uncertainties are not taken into account during the fitting, though future versions of the algorithm will take into account Poisson distribution uncertainties to minimize the chi-squared. Care will be needed in estimating the variance λ of the data in the large regions where few, if any, beta-gamma coincidence counts are recorded. A constant background component may be added to improve the algorithm's fit.

In the longer term, we hope to develop a calibration routine for the isomer peaks to automate the process for each detector, and to develop a rigorous test of the peak-fitting program against the ROI method, SDAT method, and other contemporary proposals for radioxenon measurement.

Compliance with ethical standards

Conflict of interest The authors declares that have no conflict of interest.

References

1. Bowyer T, Abel K, Hubbard C, Panisko M, Reeder P, Thompson R, Warner R (1999) Field testing of collection and measurement of radioxenon for the comprehensive test ban treaty. J Radioanal Nucl Chem 240(1):109–122
2. Bowyer TW, Schlosser C, Abel KH, Auer M, Hayes JC, Heimbigner TR, McIntyre JI, Panisko ME, Reeder PL, Satorius H et al (2002) Detection and analysis of xenon isotopes for the comprehensive nuclear-test-ban treaty international monitoring system. J Environ Radioact 59(2):139–151
3. Bowyer TW, McIntyre JI, Reeder PL (1999) High sensitivity detection of Xe isotopes via beta-gamma coincidence counting. Technical report, Pacific Northwest National Laboratory
4. McIntyre JI, Schrom BT, Cooper MW, Prinke AM, Suckow TJ, Ringbom A, Warren GA (2016) A program to generate simulated radioxenon beta-gamma data for concentration verification and validation and training exercises. J Radioanal Nucl Chem 307(3):2381–2387
5. Cooper MW, Auer M, Bowyer TW, Casey LA, Elmgren K, Ely JH, Foxe MP, Gheddou A, Gohla H, Hayes JC et al (2019) Radioxenon net count calculations revisited. J Radioanal Nucl Chem 321(2):369–382
6. Maceira M, Blom PS, MacCarthy JK, Marceillo OE, Euler GG, Begnaud ML, Ford SR, Pasyanos ME, Orris GJ, Foxo MP, Arrowsmith SJ, Merchant BJ, Slinkard ME (June 2017) Trends in nuclear explosion monitoring research & development. Technical Report LA-UR-17-21274, Los Alamos National Laboratory, Los Alamos NM. https://doi.org/10.2172/1355758
7. Haas DA, Biegalski SRF, Biegalski KMF (2009) SDAT: analysis of ^{131m}Xe with ^{133}Xe interference. J Radioanal Nucl Chem 282:715–719
8. Cooper MW, Auer M, Bowyer TW, Casey LA, Elmgren K, Elmgren JH, Foxe MP, Gheddou A, Gohla H, Hayes JC, Johnson CM, Kalinowski M, Klingberg FJ, Liu B, Mayer MF, McIntyre JI, Plenteda R, Popov V, Matthias Z (2018) Radioxenon net count calculation refinements. Technical report, Pacific Northwest National Lab, Richland, WA
9. Cooper MW, Bowyer TW, Hayes JC, Heimbigner TR, Hubbard CW, McIntyre JI, Schrom BT (2008) Spectral analysis of radioxenon. Technical report, Pacific Northwest National Lab, Richland WA
10. Biegalski SR, Biegalski KM, Haas DA (2007) Testing the spectral deconvolution algorithm tool (SDAT) with Xe spectra. Technical report, Texas University at Austin
11. Biegalski S, Flory A, Haas D, Ely J, Cooper M (2013) Sdat implementation for the analysis of radioxenon β-γ coincidence spectra. J Radioanal Nucl Chem 296(1):471–476
12. Biegalski S, Flory AE, Schrom BT, Ely JH, Haas DA, Bowyer TW, Hayes JC (2011) Implementing the standard spectrum method for analysis of β-γ coincidence spectra. In: Proceedings of the 2011 monitoring research review: ground-based nuclear explosion monitoring technologies, pp 607–615
13. Biegalski KF, Biegalski S, Haas D (2008) Performance evaluation of spectral deconvolution analysis tool (sdat) software used for nuclear explosion radionuclide measurements. J Radioanal Nucl Chem 276(2):407–413
14. Rouaud M (July 2013) Probability, Statistics, and Estimation. Incertitudes
15. Krane KS (1988) Introductory nuclear physics. Wiley, Hoboken

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.