Characterization of highly multiplexed monolithic PET / gamma camera detector modules

Paper

Pierce, L. A.; Pedemonte, Stefano ; Dewitt, Sharon; Macdonald, L.; Hunter, W. C. J.; Leemput, Van; Miyaoka, R.

Published in:
Physics in Medicine and Biology

Link to article, DOI:
10.1088/1361-6560/aab380

Publication date:
2018

Document Version
Peer reviewed version

HHS Public Access

Author manuscript

Phys Med Biol. Author manuscript; available in PMC 2019 March 29.

Published in final edited form as:

Phys Med Biol. ; 63(7): 075017. doi:10.1088/1361-6560/aab380.

Characterization of highly multiplexed monolithic PET/gamma camera detector modules

L A Pierce[1], **S Pedemonte**[2], **D DeWitt**[1], **L MacDonald**[1], **W C J Hunter**[1], **K Van Leemput**[2,3], and **R Miyaoka**[1]

[1]Imaging Research Laboratory, Department of Radiology, University of Washington, 1959 NE Pacific St., Seattle WA, USA

[2]Martinos Center for Biomedical Imaging, MGH, Harvard Medical School, Boston, MA, USA

[3]Department of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark

Abstract

PET detectors use signal multiplexing to reduce the total number of electronics channels needed to cover a given area. Using measured thin-beam calibration data, we tested a principal component based multiplexing scheme for scintillation detectors. The highly-multiplexed detector signal is no longer amenable to standard calibration methodologies. In this study we report results of a prototype multiplexing circuit, and present a new method for calibrating the detector module with multiplexed data.

Methods—A 50 × 5 × 10 mm^3 LYSO scintillation crystal was affixed to a position-sensitive photomultiplier tube with 8 × 8 position-outputs and one channel that is the sum of the other 64. The 65-channel signal was multiplexed in a resistive circuit, with 65:5 or 65:7 multiplexing. A 0.9 mm beam of 511 keV photons was scanned across the face of the crystal in a 1.52 mm grid pattern in order to characterize the detector response. New methods are developed to reject scattered events and perform depth-estimation to characterize the detector response of the calibration data. Photon interaction position estimation of the testing data was performed using a Gaussian Maximum Likelihood estimator and the resolution and scatter-rejection capabilities of the detector were analyzed.

Results—We found that using a 7-channel multiplexing scheme (65:7 compression ratio) with 1.67 mm depth bins had the best performance with a beam-contour of 1.2 mm FWHM (from the 0.9 mm beam) near the center of the crystal and 1.9 mm FWHM near the edge of the crystal. The positioned events followed the expected Beer-Lambert depth distribution. The proposed calibration and positioning method exhibited a scattered photon rejection rate that was a 55% improvement over the summed signal energy-windowing method.

1. Introduction

In scintillation-based gamma ray imaging systems, such as positron emission tomography (PET) detector blocks and general nuclear medicine gamma cameras, high energy photons are generally detected by interaction with a scintillation crystal that generates a burst of lower-energy photons that are, in turn, detected by a photosensor, generally a silicon

photomultiplier (SiPM) or a photomultiplier tube (PMT). Using a position-sensitive photosensor, the output signal can be analyzed and the position of the original photon interaction can be estimated. In this work we focus on a monolithic crystal detector module for PET imaging.

The Continuous Miniature Crystal Element (cMiCE) PET detector modules currently under development at the University of Washington utilize a 50 × 50 × 10 mm^3 continuous block of LYSO scintillation crystal coupled to a 65-channel multi-anode photomultiplier tube (MA-PMT) (Miyaoka et al. 2011)(Miyaoka et al. 2007). The MA-PMT has an 8×8 array of photosensors as well as a 65^{th} output channel that is the sum of the other 64 outputs. The estimation of interacting photon positions is performed using Gaussian maximum likelihood (ML) (Ling et al. 2008) (Ling et al. 2006).

The Gaussian ML position estimator requires a lookup table of the mean and variance of the PMT output signal for every (*x, y, z*) photon position within the crystal (segmented into discrete (*x, y, z*)-bins). The mean and variance are obtained by scanning a thin beam of 511 keV photons across the face of the scintillation crystal and recording the photomultiplier output. The calibration beam is generated using a ^{22}Na point source, a coincidence crystal, and a Tungsten shield.

The 511 keV calibration-beam data can be contaminated in several ways, including annihilation photons that have scattered off of the Tungsten collimator, high-energy (1,274 keV) gammas from the ^{22}Na point source that passed through or scattered off of the collimator, photons that scattered within the scintillation crystal, or other types of random events. Identifying and rejecting scattered and random events in the calibration data is critical to the generation of accurate lookup tables for the Gaussian ML photon positioning algorithm. Scatter rejection for calibration data was previously performed on cMiCE detectors using a two-step process ((Ling et al. 2006), (Miyaoka et al. 2011), (Pierce et al. 2014)). An initial scatter-rejection was performed via an Anger position estimate of the 8 × 8 position-sensitive PMT outputs. The Anger position estimate of each event was then compared to the Anger position estimate of the mean signal over all events and a median-distance threshold was used to reject events whose Anger position was too far from the mean position. A second filtering was performed by using the 65^{th} summed-channel (a common dynode, with output related to the sum of the 8×8 position-sensitive channels) of the PMT to measure the energy spectrum for all photons from a given beam position. This spectrum was used to define an energy window around the photopeak. Each event was then screened against this summed-channel energy window as the second round of scatter-rejection.

After scatter rejection has been applied to the calibration data, only the (*x, y*)-position of each event from a given calibration beam is known. The *z*-direction, being the depth-of-interaction of the photon is randomly distributed according to the Beer-Lambert distribution. The expected output signal for an interacting photon varies as a function of the depth of interaction within the crystal, requiring that the detected events (after scatter-rejection) be sorted according to depth in order to accurately localize the calibration photon interactions and create 3-dimensional lookup tables for the Gaussian ML position estimator.

Calibration photon depth estimation was previously performed by clustering the outputs for a given beam position according to the strength of the maximum channel from the 8 × 8 position-channel outputs (not using the summed-channel) (Ling et al. 2007). That method made use of the fact that the photons near the photomultiplier will have strongly localized signals and those further away will have their signal spread out more, resulting in a relatively lower peak channel signal. In later works, the depth of interaction for a given beam was performed by fitting a Gaussian, Cauchy, parametric, or Lorentzian curve to each 8×8-channel event and depth-sorting the events according to the fitting function parameters (Ling et al. 2008), (Pierce et al. 2014). These methods require knowledge of the shape of the 8 × 8-channel signal in order to fit the depth-estimation function to the signal output. These scatter rejection and depth estimation methods rely on having data from all 65 output channels from the PMT.

In Pierce et al. (2014), a data-driven multiplexing scheme was proposed that utilized the principal components of the expected output signal (PCA-MUX). In that work, the realization was made that general signal reconstruction for photon positioning was not necessary. This is because each incoming photon will result in PMT outputs that have a specific output “shape”, and not all output shapes are valid. Thus, the expected output data lie near a lower-dimensional manifold within the 64-dimensional output vector space. A dataset was collected by scanning a collimated beam across the cMiCE detector face. By taking the principal components of this dataset, each photon interaction output signal could be described as a weighted sum of only the most significant principle components. This resulted in reasonable signal reconstruction with as few as 8 principal component output channels.

However, the highly multiplexed PCA-MUX data do not preserve the information needed for the scatter rejection and depth estimation techniques mentioned above. In Pierce et al. (2014), the 65-channel data was acquired and the multiplexing of the 8 ×8 position channels was performed via simulation. Scatter rejection and depth estimation were performed using all 65 channels of the data prior to virtual multiplexing.

In a follow-up work (Hunter et al. 2017), the PCA-MUX scheme from Pierce et al. (2014) was encoded in hardware in a resistive circuit and attached to a cMiCE detector module. In that work, a collimated beam was scanned across the detector prior to the attachment of the multiplexing circuit, resulting in a 65-channel output being acquired. Depth estimation of photons and scatter rejection were performed on the non-multiplexed dataset. Scatter rejection was performed by defining an energy window on the summed-signal channel for each beam position. Depth estimation was performed by fitting a Lorentzian distribution to each photon signal and clustering events according to the fitted Lorentian parameters. Each of the 64 inputs of the PCA-MUX circuit was probed with a known pulse and the outputs recorded. These pulse outputs were used to simulate the effects of the multiplexing circuit on the PMT output. The filtered and depth-separated data were then multiplexed in simulation and Gaussian ML lookup tables were created for positioning. The PCA-MUX circuit was then attached to the cMiCE detector module and a second set of calibration beam data was collected. This hardware-multiplexed dataset was then positioned using the lookup tables generated by simulated multiplexing and detector performance was evaluated.

Both Pierce et al. (2014) and Hunter et al. (2017) suffer from the same hindrance in that the 64-channel non-multiplexed data needed to be collected prior to affixing the multiplexing circuit in order to create accurate Gaussian ML lookup tables as well as for scatter rejection and depth estimation. In order to use the traditional scatter rejection and depth-estimation calibration methods with the PCA-MUX method, the 65-channel (non-multiplexed) data would need to be acquired at calibration time in order to characterize the multiplexed detector response. This would require a 65-channel “calibration mode” addition to the multiplexing circuitry, adding to the cost and complexity of the method.

For characterization of the detector module, it is beneficial to create new scatter rejection and depth-estimation algorithms that can work directly on the highly-multiplexed PCA-MUX signal output obtained from the calibration beam.

Later work used machine learning and dimensionality reduction techniques to perform the depth separation step for calibration of monolithic crystal detectors (Pedemonte et al. 2017). In that work, multiplexing was not considered and only non-multiplexed signals were analyzed. A collimated beam was scanned across the face of cMiCE detector and each beam was analyzed using the Locally Linear Embedding (LLE) algorithm (Roweis & Saul 2000) to determine the depth at which each photon interacted. The realization was made that despite the 65-dimensional output from the detector, the data from each calibration beam is intrinsically one-dimensional within the output data space, with depth being the only dimension. The transformation from (x, y, z) 511 keV photon interaction position to 65-channel PMT output is non-linear, but mostly continuous, implying that the output data lie along a 1-dimensional manifold embedded in the 65-dimensional data space. The LLE algorithm was able to determine where on that manifold each data point lie and thus determine the depth of interaction.

In this paper, we combine the methods of Pierce et al. (2014), Hunter et al. (2017), and Pedemonte et al. (2017). We report here the hardware implementation of the 5-and 7-channel PCA-MUX multiplexing scheme proposed in Pierce et al. (2014) and implemented in Hunter et al. (2017). For characterization of the multiplexed data, we follow the reasoning of Pedemonte et al. (2017) and utilize the dimensionality reduction technique of Principal Components (Pearson 1901). We report the creation of a new algorithm, called PC3 (because the method of principal components is used three times in the calibration workflow), for the highly-multiplexed detector data. The PC3 algorithm works directly on the multiplexed dataset and eliminates the need to collect non-multiplexed data for detector calibration. Both scatter rejection and photon depth estimation are performed within the PC3 algorithm to build accurate Gaussian ML lookup tables.

For this work, we construct the PCA-MUX multiplexing scheme using a resistive circuit with 65:7-channel and 65:5-channel outputs. We test the performance characteristics of the cMiCE detector module using the new PCA-MUX multiplexing and calibration methodologies with the PC3 algorithm.

2. Methods

2.1. Design and Creation of the Multiplexing Circuit

In order to determine the weights (relative output) of our PCA-MUX multiplexing circuit, we created a simple simulation for the output of interacting photons. A 50 ×50 ×10 mm^3 crystal was modeled. Only single-interaction photons were simulated, all with equal energy, and the output from each of the 8 × 8 position channels was recorded as the solid angle from the interacting photon position to each of the 64 detector elements (the 65^{th} summed-channel was not used). The entrance face of the crystal was modeled as a perfect specular reflector, the edges perfectly blackened as shown in Figure 1. Isotropic photon point sources were simulated on a cubic grid within the crystal with a 1 mm spacing in all three dimensions. No noise was modeled in this simulation. (i.e. This simulation was a geometric calculation only.)

The principal components of the 8 × 8-channel simulated dataset were then computed, and these were then subsequently scaled and rounded so that only seven values were present in the principal component vectors: {−3, −2, −1, 0, 1, 2, 3}, as in Pierce et al. (2014). The 16 most significant principal components from this dataset were tested via virtual multiplexing of the previously collected cMiCE dataset from Pierce et al. (2014) to determine which combination of output channels performed the best (visual inspection of the half-max contours, FWHM of the half-max contours, minimal positioning bias, performance at the edge of the crystal, and agreement with the Lorentzian-fit depth estimation were all considered in this choice). A set of five principal components was chosen for the hardware implementation of PCA-MUX (numbers 1, 2, 4, 8, and 9 in order of descending principal component singular value).

The same geometric simulation was run again using only photon interaction points within 3 mm of the edge of the crystal (approximately one-half the width of one detector element). This produced 64 principal components, which were scaled and rounded to 7 discrete values (as above), and two of the “edge-specific multiplexing channels” were chosen to accompany the 5 channels previously chosen (numbers 5 and 6 in order of descending principal component singular value). This choice was based on the ability of the edge-specific multiplexing channels to enhance the performance of the detector near the edge of the crystal when used in addition to the 5 channels previously chosen.

These seven sets of weights (5 principal components plus 2 edge-specific components) were then coded into a resistive circuit that was connected to the cMiCE detector module (Hunter et al. 2017). Scaling and rounding the real-valued weights of the principal components ensured that only 3 resistor values needed to be used in the construction of the circuit. Negative values were achieved via voltage inverters.

Once the PCA-MUX circuit was constructed, each of the 8 × 8 position-channels of the circuit input was probed with a known voltage (similar to the expected response from an incoming photon interaction from the PMT) and the output of the seven multiplexed channels were recorded. The Euclidean norms of each of these 7-dimensional vectors from

the test pulses was recorded as $\{W_i\}_{i=1}^{7}$, the pulse weights for later signal normalization.

The resulting pulse weights are shown in Figure 2.

2.2. Collection of Calibration Data

The PCA-MUX resistive circuit was connected to a cMiCE detector module (Miyaoka et al. 2011), which consists of a monolithic slab of LYSO scintillation crystal (Crystal Photonics Incorporated) measuring 50 × 50 × 10 mm^3 affixed to an 65-channel MA-PMT (H12700A, Hamamatsu Photonics K.K., Japan). The entrance face of the crystal was covered with 0.015 inch Teflon (PTFE) sheet and then a 3M® ESR mirror-film reflector, and the sides were painted with flat black latex paint. The crystal is directly coupled to the MA-PMT window with BC-630 silicon optical grease by Saint Gobain Crystals®. Note that the 65^{th} (common-dynode) channel of the H12700A MA-PMT is not used in our calibration or positioning algorithms, but is used when we compare the proposed method to traditional energy windowing for scatter rejection as outlined in Section 2.4. Also note that this configuration is not necessary for the proposed calibration methodology, which we believe will work for any choice of monolithic crystal size, composition, or reflective or blackened surfaces.

A thinly-collimated beam was created using a 0.5 mm diameter ^{22}Na point source (Isotope Products Laboratories, Valencia, CA.), collimated using a 4 × 4 × 20 mm^3 coincidence crystal placed 11 cm beyond the point source. The point source was 10 mm from the face of the cMiCE crystal. A tungsten collimator with a 2.0 mm diameter hole was used to keep the point source from flooding the rest of the detector to mitigate event pile up. This resulted in a beam that was 0.93 mm FWHM in diameter at the entrance face and 0.98 mm FWHM in diameter at the exit face of the cMiCE crystal. This beam was scanned on a 1.52-mm grid across the face of the detector module, resulting in a 33-by-33 set of (x,y)-beam positions. Our data acquisition electronics triggered on the common-dynode signal (65^{th} channel) and collected the multiplexed calibration-beam data for an ensemble of 25,000 events at each beam position. The signal for each output channel from the multiplexing circuit was then normalized by dividing the recorded pulse weight W_i, determined in Section 2.1. (This is done to ensure that the maximum likelihood estimation is not dominated by any channel with larger voltage output.) The pulse weight normalized data from each beam was then split in half: one half for calibration of the detector, the other half for testing the performance of the module.

Full details of the creation of the multiplexing circuit and the collection of the calibration beam dataset are in Hunter et al. (2017).

2.3. The PC3 Calibration Method

We perform several sets of experiments with the dataset from the PCA-MUX circuit using two multiplexing schemes: one using seven output channels (5 full principal component channels plus the two edge-specific channels), and also using only 5 output channels (using the same calibration dataset, but ignoring the output from the two edge-specific channels). Throughout this paper, we refer to the number of output channels as ch (with ch = 5 or 7). Note that we do not use the summed-channel for the proposed PC3 calibration or positioning the testing data.

We also vary the number of positioning depth bins between 6, 7, and 10 depth bins. Each depth bin has a fixed physical size (in millimeters: 1.67, 1.43, and 1.00 mm respectively), resulting in depth bins that do not have the same expected number of counts. We refer to the number of depth bins as D in what follows.

Overall, this yields six total experiments that we report on: {2 multiplexing schemes} × {3 depth-bin widths}.

The PC3 calibration algorithm is predicated on the idea that the calibration beam data lies on a one-dimensional path (manifold) within the higher-dimensional space of output signals. The most significant principal component should indicate the “depth direction” within the data space. However, in our dataset, the largest data variance for each calibration beam is actually due to scattered and random photon interactions.

Our proposed method computes the principal components of all beam data for initial density-threshold-based scatter rejection. This allows us to accurately find the depth-direction as the most significant principal component of the remaining data. This is illustrated in Figure 4. Once the depth-direction is determined, a density-based threshold filtering is applied to the data in each dimension orthogonal to the depth-direction.

For each beam position, the principal components of the calibration data were computed, we call these vectors $\{PC_{1,i}\}_{i=1}^{ch}$. The output data events were then transformed into the principal component coordinates via change-of-basis. The transformed channel values were then histogrammed (a one-dimensional histogram for each transformed channel). This gave a density profile of events according to the principal component coordinate values. Upper and lower thresholds were defined as a fraction of the maximum density as shown in Figure 5(a). Events that fell outside of this window were rejected as scattered events.

The principal components of the data that passed the initial scatter rejection were then computed resulting in principal component vectors $\{PC_{2,i}\}_{i=1}^{ch}$. The density of values of the most significant principal component, $PC_{2,1}$, was analyzed and upper and lower thresholds were determined according to this density profile, as shown in Figure 5(b). Thresholds were computed as a fraction of the maximum density. Events that fell below the lower threshold were rejected as low-energy scatter, events that fell above the upper threshold were rejected as high-energy ^{22}Na gamma events. Note that the initial scatter rejection was performed only to ensure alignment of the $PC_{2,1}$ vector as shown in Figure 4(c).

All of the calibration data events were then projected onto the data subspace orthogonal to the $PC_{2,1}$ vector, resulting in one less dimension spanned by the transformed data ($ch-1$ dimensions). The principal components of the projected data were then computed. The data that passed the initial $\{PC_{1,i}\}_{i=1}^{ch}$ filtering were then transformed into the new principal component coordinates and a third round of density-based scatter rejection was performed along each of those $ch-1$ dimensions (Figure 5(c)). The thresholds previously defined along the $PC_{2,1}$ dimension were again used along that dimension, resulting in a ch-dimensional bounding box.

Those events that lie within this *ch*-dimensional bounding box were considered qualified events, and the inner product with the event data and the $PC_{2,1}$ vector was used as a surrogate for the depth of interaction of that event (a higher inner product value implying a deeper depth). Depth bins were defined according to the expected fraction of interactions per bin, according to the Beer-Lambert distribution for 511 keV events in LYSO crystal with attenuation coefficient = 0.087 mm^{-1}. Once each qualified event was assigned a depth bin, the mean and standard deviation of the multiplexed signal was recorded for each depth bin.

In order to perform a faster 2-stage Gaussian maximum likelihood positioning, two sets of lookup tables were recorded. The first set, which we call $\mu_{initial}$ and $\sigma_{initial}$, use 3 depth bins resulting in $33x \times 33 \times 3 \times ch$ lookup tables. The second set of lookup tables, μ_{final} and σ_{final}, uses D depth bins and is interpolated in the x and y-directions to 0.30 mm bins using 2D spline interpolation. Each of the four lookup tables is smoothed using 3D Gaussian smoothing on each output channel.

When computing the $\mu_{initial}$ and $\sigma_{initial}$ lookup table values, the Euclidean norm of the events in each (x, y, z) bin is determined. Upper and lower thresholds are defined according to this density profile and recorded as $E_{x,y,z,lo}$ and $E_{x,y,z,hi}$.

Each of the principal component scatter rejections requires one or more parameters to define the threshold for scatter rejection, as does the Euclidean norm window. The beams in the central region of the crystal (more than 3.5 mm from the edge) were treated equally. The parameters were gradually tightened outside the central region so that as the beams get closer to the edge of the crystal, more photons are rejected as scatter (the fifth from the edge of the crystal being the first beams differing from the central region). Specific details of these parameters can be found in the matlab source code, which is available upon request from the corresponding author.

2.4. Photon Positioning and Performance Metrics

For each event in the pulse weight normalized testing dataset, the likelihood values are computed according to the $\mu_{initial}$ and $\sigma_{initial}$ lookup tables and the argmin over (x, y, z) is chosen as the initial position estimate for the photon. The Euclidean norm of that event signal is then compared to the $E_{x,y,z,lo}$ and $E_{x,y,z,hi}$ Euclidean norm thresholds for the initial (x, y, z) position estimate. The event is rejected as scatter if the Euclidean norm falls outside this window. If the event lies within the Euclidean norm window (an “energy qualified event”), a second set of likelihood values is computed in the local neighborhood in a 3×3 mm^2 square region centered on the initial position estimate (over all depths), but using the interpolated lookup tables μ_{final} and σ_{final} and its position recorded.

The energy qualified and positioned test events are then spatially histogrammed into a 0.30 × 0.30 × (10/D) mm^3 array. Half-max contours are created by summing the positioning histogram over all depths and the full width at half maximum (FWHM) was computed for each test beam. The fraction of events positioned per depth bin is also recorded for each beam position and compared to the expected Beer-Lambert distribution.

The mean over all FWHM were computed for three regions: the center of the crystal, the edges of the crystal, and the corners of the crystal. The edge and corner regions are defined by the outermost four beam rows and columns; the central region is the remainder. In the central region, the FWHM is measured in the *x* and *y*-directions. In the edge and corner regions, the FWHM is measured in the 'radial' and 'tangential' directions (perpendicular or parallel to the nearest edge or corner, as in Pierce et al. (2014)).

2.5. Analysis of Scattered Photon Rejection Capability

The ability of a detector module to recognize scattered events is an important performance consideration. Often, a photon emitted from a position within the scanner field of view will interact via Compton scattering prior to entering the detector module. This scattered photon will deflect from its original path, leading to an incorrect entry into the data set (a sinogram or other data array) for the scanner. This degrades scanner performance in both image resolution and quantitative accuracy.

Fortunately, Compton scattering causes the photon to lose energy prior to entering the detector. In order to identify scattered photons, it is common to estimate the energy of each photon that interacts within the detector. In this way, photons that have estimated energies that are too low can be rejected as scattered events. (An upper energy threshold can also be defined to eliminate high-energy pile-up events, random cosmic rays, etc.).

In general, a given scan protocol will utilize a single photon energy to obtain tomographic data. In this case, the detector need only be trained to accept or reject photons for the anticipated monochromatic photon energy. If one were to create a theoretically accurate histogram of the detected energies of all photon events from a monochromatic source, a single energy spike, called a "photopeak", would be visible, which would be separated from a broad polychromatic range of lower-energy scattered events that entered the detector. Theoretically, a monochromatic photopeak should be a delta function, consisting of a peak at the single output voltage value that represents the anticipated photon energy.

Both the photopeak and scattered event histogram will be blurred in a measured histogram due to imperfections in the measurement process. Illustrations of energy histograms are shown in Figure 10, a measured energy histogram is shown in Figure 11. The width of the measured (blurred) photopeak is often used to compute the "energy resolution" of the detector. This energy resolution is used to estimate the energy discrimination capabilities of the detector and thus the detector's ability to reject scattered events.

A common method to estimate photon energy is by using the sum of all output signals from the photosensor. The summed output voltage is histogrammed and the photopeak is determined by examining the shape of this distribution (clearly seen as a large spike at the upper end of the histogram as seen in Figures 10 and 11). An "energy window" is defined around the photopeak and detected events that have summed-channel voltages outside of this energy window will be rejected as scattered photon events.

Unfortunately, in monolithic crystals, the theoretical photopeak for a given (*x, y*)-position is no longer a delta function at a single output voltage. Photons that interact near the

photosensor will result in more photons detected by the (nearby) photosensor compared to a photon that interacts far away from the photosensor (near the entrance surface of the crystal). This means that two photons that deposit the same energy at different distances from the photosensor will have two different output voltages. For this reason, the standard metric of energy resolution as the width of the photopeak is not appropriate for the cMiCE detector module calibration. We will propose a new metric to assess the scatter-rejection capabilities of the PC3 method for monolithic crystals.

A confounding factor affecting the scatter-rejection abilities of the monolithic detector module is spatially variant energy output. This can be due to inhomogeneities and defects within the crystal, as well as voltage bias variations throughout the photosensor. The result is that the output voltage photopeak position can vary throughout the detector. This is illustrated for two hypothetical detector positions in Figure 10.

This was overcome in previous work (Pierce et al. 2014) by computing an energy window for each calibration beam position in the detector. During an imaging task, each incoming photon undergoes an initial Gaussian ML position estimate, then the summed-signal output for that photon is compared to the energy window recorded for that position in the detector. We refer to this method as the “initial-positioning scatter rejection method”.

As an illustration, consider a detector that is calibrated to 400 keV photons using a scanned thin beam of photons. For two positions in the detector module (x_0, y_0) and (x_1, y_1), Suppose that the energy histogram for the calibration beam at position (x_0, y_0) shows a photopeak at 8 volts, while the photopeak for position (x_1, y_1) is at 5 volts as shown in Figure 10.

During an imaging task, suppose that a radiotracer in the field of view emits a 400 keV photon that experiences a Compton scattering prior to entering the detector. Suppose that this photon enters the detector at position (x_0, y_0) with 250 keV of energy, resulting in a summed-signal output of 5 volts.

Due to the lower voltage output of our scattered photon, the photosensor signal is more likely to match the lower voltage values at position (x_1, y_1) in the Gaussian ML lookup table. Thus, the initial Gaussian ML positioning estimate for the scattered 250 keV photon will be at (x_1, y_1) even though the photon is physically at position (x_0, y_0) in the detector. Furthermore, the 250 keV photon’s lower output voltage will be within the energy window at (x_1, y_1), resulting in the scattered 250 keV photon at position (x_0, y_0) being recorded as a 400 keV photon at position (x_1, y_1).

If our example photon had an initial position estimate that matched its physical position, we could use the energy window for its true position and see that the photon should have been rejected as scattered. Fortunately for us, each photon from our testing dataset has known (x, y)-position and we can use the correct energy window to determine if each photon should have been rejected as scattered according to the energy window of the photon’s physical position.

We introduce the “Locally Qualified Acceptance Fraction” (LQAF) as a metric to assess the scatter rejection capabilities of a detector. For a given beam position, we use the calibration

dataset to histogram the 65^{th} summed-channel from the PMT and define an energy window. For the testing data, we record the final fate of each event: was it rejected as scattered, or was it accepted and positioned (possibly at an incorrect position within the crystal). The LQAF is defined as the fraction of events that were accepted (not rejected as scatter) and that also fell within the energy window for the beam that generated that event (the energy window of the true physical event position). An LQAF of 1.0 (100%) would mean that no accepted photons fell outside of the local energy window (a perfect scenario). An LQAF of 0.0 (0%) would indicate that every accepted photon fell outside of the correct energy window (all photons were mis-positioned and accepted by energy windows away from the true beam position).

For our dataset, we use the 65^{th} summed channel as a surrogate for energy, create a histogram of the output voltages from that channel for the calibration data, and define an energy window based on that histogram. This energy window is defined by a lower threshold of 5/6 of the photopeak (halfway to the Compton edge) and an upper threshold that is 1.5 times the half-width at tenth-max above the photopeak (the same as in Pierce et al. (2014)).

We then position our testing dataset according to the PC3 method and histogram the summed-channel energies of all PC3-accepted photon events. The LQAF is computed as the fraction of events accepted by PC3 that also lie within the local energy window. We then repeat the LQAF computation using the “initial-positioning scatter rejection method” mentioned above and previously used in Pierce et al. (2014). The LQAF is computed for both scatter-rejection techniques over all $33 \times 33 = 1,089$ beam positions.

3. Results

The mean FWHM for each region are reported in Table 1. Half-max contours for beam positions are shown in Figures 6 and 7. Comparison of the 7- and 5-channel contour maps shows that the addition of the two edge-specific channels does improve the performance near the edges of the detector. In the 5-channel experiment shown in Figure 7, the half-max contours are crowded together near the edge, while the contours are generally better separated in the 7-channel experiment shown in Figure 6. The contours along the top and bottom edges show less bias in the 7-channel experiment. The contours along the left edge exhibited large biases away from the edge in every experiment.

We show the fraction of events positioned per depth for the six experiments in Figures 8 and 9. We note that altering the number of depth bins did not have an appreciable effect on the half-max contours or the FWHM measurements for either the 7-channel or 5-channel experiments, with the exception of a few local biases as evidenced in Figures 6 and 7.

The distribution of the number of events positioned per depth bin are shown in Figures 8 and 9. We can see that the distributions for each beam are closely related. The results for using 6 depth bins show a distribution that is close to the theoretic Beer-Lambert distribution and becomes more erratic as we increase the number of depth bins. We note that altering the number of output channels did not have an appreciable effect on the distribution of events per depth.

The scatter-rejection histograms of a single beam position are shown in Figure 11. Other beam positions have similar histograms except for a few beam positions near the edges and corners. The summed-channel accepted-event histograms show that scattered events (likely scattered prior to entering the crystal and lower energy than 511 keV) from this beam were positioned elsewhere in the crystal and accepted by the summed-channel energy window at the estimated position, as evidenced by the low-end distribution of the black line. However, the sharply-peaked red line shows that scattered events from this beam that were positioned elsewhere were generally not accepted by the PC3 method.

Figure 12 shows the results of the “locally qualified acceptance fraction” (LQAF) for each beam position using the 7-channel, 6-depth experiment. It is clear that the PC3 method shows a more homogeneous and generally higher fraction of correctly-accepted events within the local energy window. The LQAF becomes higher near the edges of the crystal for the PC3 method. The region of increased LQAF coincides with the region in which we altered the parameters of the PC3 algorithm to enforce stricter scatter-rejection bounds. We did not test for correlation or causation for the increased LQAF values near the edge. The mean LQAF over all beam positions for the summed-channel method is 58.1%. The mean LQAF for the PC3 method is 90.1%, a 55% improvement. For the 5-channel experiments, the mean LQAF was 61.5% using the summed channel for energy discrimination, and 89.4% for the PC3 method, a 45% improvement. The colormaps for other experiments are very similar to those for the 7-channel, 6-depth experiments shown in Figure 12. There was no appreciable difference when the number of depth bins is changed.

4. Discussion

During the process of acquiring the multiplexed data, unfortunate laboratory circumstances resulted in our not being able to obtain slant-beam or side-beam illumination of the detector module, nor a calibration beam profile.

To estimate the calibration beam diameter, we computed the analytic beam profile of a true point source with the described coincidence setup. We convolved this with the profile of the 0.5 mm point source (including positron range) to obtain the estimated 0.93 mm FWHM calibration beam diameter (0.98 mm FWHM at the rear of the crystal). The collimator was used to shield the rest of the monolithic detector from being flooded, and was not used to shape the calibration beam and was not considered in our computation. It is possible to de-convolve the measured data with the analytically-estimated beam profile. However, without a measured beam profile, this would be speculative and could yield exaggerated results.

Our FWHM resolution measurements do not represent the intrinsic resolution of the detector, but include the distribution of the 0.9 mm diameter beam profile. Yet, the FHWM measurements presented in this paper are similar to those presented in the previous work (Pierce et al. 2014). However, the FWHM measurements for that previous work were determined from a calibration beam that was 0.52 mm in diameter (stepped in a 1.013 mm grid), while the current work uses a 0.9 mm diameter beam (stepped in a 1.52 mm grid). This implies that the FWHM from the new method is significantly better than those presented in Pierce et al. (2014).

Another limitation of not having a beam profile is that we cannot estimate the true number of scattered events from the calibration beam. Still, we believe that the locally qualified acceptance fraction defined in this paper illustrates a that the proposed PC3 method is a significant improvement over the summed-signal energy-windowing technique (55% improvement for 7-channel experiments and 45% improvement for 5-channel experiments). However, the LQAF metric is likely appropriate for the present work only.

Not having slant- or side-illumination of the detector means that we cannot accurately estimate the depth-resolution of the detector. Even so, we see that the distribution of events roughly follows a Beer-Lambert distribution for 6 depth bins (1.67 mm bins) and degrades as we increase the number of bins. This is true for both the 7-channel and 5-channel experiments. While this cannot tell us the depth resolution, it gives some indication of the depth-positioning capabilities of the detector using the PC3 method. The “bump” in the mid-deep bins will be the subject of future investigation, once side-illumination data can be obtained.

In section 2.1 we briefly discussed how we chose the principle components of the simulated data and how those principle components were encoded in the multiplexing circuit. We tested (via virtual multiplexing of acquired data) many combinations of the first 16 principle components as well as the inclusion of many types of edge-specific channels. Our final choice consisted of 5 principle components and 2 edge-specific channels. The choice of which of the principle components and edge-specific configurations to include was partly based on quantitative measures like the FWHM of the positioning histogram profiles, but also on qualitative observations as well. As an example, Figure 13 shows half-max contours for the 7-channel experiment presented throughout this paper, but with channel 3 (principle component 4 of the simulated dataset) omitted (thus only using 6 channels). While the contour map gives a reasonable quantitative measure of 1.3 mm FWHM, unacceptable biases are seen in the contour map. Many of the beams are projected to a single point (all photons perfectly positioned to the beam center, resulting in an unrealistic 0.0 mm FWHM) while other beams are positioned far away from the true position.

It was found that the 4^{th} principle component was necessary for reasonable positioning, as were principle components 1 and 2. However, we found that principle components 3 and 5 were not necessary and that replacing them with components 6 and 9 led to better overall performance. It was observed that some of the principle components must be used in pairs, while others did not, similar to the Zernicke polynomials (Zernike 1934). It was also found that the multiplexing channels act as a pre-whitening of the data, allowing for better performance of the Gaussian ML position estimator. Principle channels were not chosen by their pre-whitening properties, but those that were chosen appear to act as better pre-whiteners of the data throughout the detector.

The methods for the choice of principe components to include in our multiplexing circuit was not exhaustive (we only examined the first 16 principle components), and we have no reason to expect that the chosen seven are optimal. However, this proof-of-concept work shows excellent resolution and energy discrimination while only using 7 output channels.

Determining the optimal channels to include in our circuit and reducing the number of output channels needed is a future direction of this work.

For both the 5- and 7-channel experiments events-per-depth distribution was similar to the theoretical Beer-Lambert distribution when using 6 or fewer depth bins. The shape of the distribution for 5 and fewer bins is very similar to that for 6 bins. When using 7 depth bins, we see that the distribution becomes more jagged and begins to deviate somewhat from the theoretical distribution. This jagged quality becomes more extreme as we increase the number of depth bins, as exhibited in the 10-depth-bin experiments shown in Figures 8 and 9. These are the only two qualitative behaviors observed in the depth distributions. We feel that presenting 6, 7, and 10-bin experiments encompasses all of the observed behavior of the depth distributions and illustrates the change in behavior between 6 and 7 depth bins (1.67 mm and 1.43 mm bin widths).

Careful inspection of the data from Figures 6(a) and 7(a) shows that the left side of the detector module suffers degradation of performance, as the half-max contours from the edge-beams were positioned away from the edge. This was true for all experiments. We believe that the calibration beam may not have been perfectly aligned and that the beam may have only been partially on the crystal in that region. This may also mean that the right side of the detector was too far onto the detector, making it easier to position events in those regions. The top and bottom edges appear symmetric, leading us to believe that good calibration beam alignment was achieved in that direction.

If this asymmetry was indeed due to mis-positioning of the crystal and calibration beam, then the boundary of the crystal (shown as blue lines in Figures 6 and 7) will be offset. This would result in a systematic bias in all contours shown in Figures 6 and 7, where each contour is biased to the left. Even if this is the case, the resolution measurements presented here will not be affected except for those positions nearest the edge.

We also see localized biased regions in the half-max contours, like the lower-right corner of the 7-channel experiments (the “blank spot” in the lower-right corner of Figure 6), and the corners of all experiments. Future work will be dedicated to improving the performance near the edges and corners of the crystal to obtain a more homogeneous response. Even with these localized biases, we see that the inclusion of the edge-specific multiplexing channels results in generally improved performance near the crystal edge and corners. This is especially apparent near the left edge: the half-max contours for the 5-channel experiment appear crowded together out to the seventh column from left, while the contours in the same region are well-separated when the edge-specific channels are included. Similar behavior is apparent along the top edge and other regions as well.

The proposed PC3 method performs better than nearly all multiplexing methodologies proposed in Pierce et al. (2014) while still using fewer output channels than those methods. All experiments in that work used the 65^{th} summed channel, while the proposed PC3 method does not. This makes the technique available to other photosensors (like silicon photomultipliers) that do not have the summed-channel capabilities of the Hamamatsu H12700A MA-PMT.

Furthermore, all experiments in Pierce et al. (2014) performed energy windowing on the testing data prior to positioning by using the 65^{th} sum-channel. This means that the known beam position and "correct" energy window was used for test data events. In practice, an incoming photon position is not known *a priori*, and an initial positioning estimate must be performed in order to compare the energy of that event with the energy window of the estimated position. This is in contrast to the current work, in which no pre-processing or filtering was performed on the testing dataset prior to positioning.

The PC3 method proposed in this work yields an uncorrected resolution of approximately 1.17 mm from the 0.93 – 0.98 mm calibration beam for the 7-channel output experiments or 1.20 mm from the 5-channel experiments (Table 1). This is comparable to the 16-channel row/column summation multiplexing result of Pierce et al. (2014), where a resolution of 1.16 – 1.25 mm FWHM was reported. However, the calibration beam in that work had a diameter of 0.52 mm, thus we expect that the PC3 results are actually better than the previous work.

The PC3 method exhibits detector resolution comparable to other methodologies, while still utilizing fewer output channels for positioning. Similar resolution results were found in other works: 1.54 mm FWHM from 0.9 mm beam (Maas et al. 2009); 1.56 mm FWHM from 0.54 mm beam (Schaart et al. 2008); 1.6 mm FWHM from 0.52 mm beam (Jorge et al. 2016); 1.45 mm FWHM from 0.5 mm beam (Borghi et al. 2016). There has been some work using thin monolithic crystals that show significant improvements in resolution: 0.60 mm FWHM resolution using a 5 mm thin monolithic crystal (Marcinkowski et al. 2016) and 0.54 mm FWHM using 2 mm thin monolithic crystal (Espana et al. 2014). In the case of thin crystals, sensitivity is sacrificed for the improved resolution, which may be desirable in some imaging applications.

The entire PC3 method (computation of lookup tables) takes about 3 minutes on a reasonable laptop computer (no parallel processing). Thus, the PC3 method can be utilized for detector calibration without the need for high-end computational equipment. The positioning of the test data takes another 40 minutes.

Due to the nature of the PCA-MUX method, similar results should be achieved using 7 output channels regardless of the number of input channels (e.g. a 12 × 12 silicon photomultiplier should still be able to have only 7 output channels with similar or better performance). However, for photosensors with more output channels, more resistive circuit values may be required to separate the PCA-MUX signals.

While the PC3 method works well for the initial calibration of a detector module, it is not feasible to use for daily quality control scans. The collimated beam scan for this dataset took approximately 18 hours (1 minute at 33 × 33 beam positions). This data acquisition also required careful alignment of the detector and a robotic gantry to move the point source.

Daily quality control scans are best performed with the detector within the scanner, as frequently removing the detector from the scanner is cumbersome and time consuming. Such scans should also be fast, so as to minimize downtime of the scanner.

One method for daily scanner calibration is an iterative method applied to flood data using the most recent lookup table as a starting point for the iterative process (Solovov et al. 2012) (Morozov et al. 2015). Another appealing approach is to use orthogonal fan-beam scans and the *k*-nearest neighbors clustering algorithm for faster calibration (van Dam et al. 2011). Methods for daily calibration of the lookup tables will be a necessary future direction of this work.

5. Conclusion

We have developed a new multiplexing scheme, called PCA-MUX, and implemented it in hardware with 7 or 5-channel output (65:7 or 65:5 compression ratios). A new algorithm, called PC3, was applied to traditional pencil-beam calibration setup. The PC3 calibration step works directly on the multiplexed data, saving the need to collect “raw” (non-multiplexed) data prior to multiplexing during the calibration phase.

The PC3 algorithm performed better than previous multiplexing and positioning methodologies with 1.2 mm FWHM contours from a 0.9 mm calibration beam. The PC3 method exhibited a scatter-rejection rate that was a 55% improvement over the previous technique of using the summed-signal for energy-windowing. The PC3 scatter-rejection also exhibits more homogeneous behavior over the crystal.

Overall, the combination PCA-MUX and PC3 are a significant improvement over previously reported multiplexing and positioning algorithms for PET detectors with monolithic crystals.

Future work will be dedicated to improving the performance at the edge and corners of the crystals, investigating more advanced clustering techniques to align the depth-direction, obtaining calibration beam profiles and side-illumination to estimate intrinsic detector resolution, and exploring the parameters for PC3 scatter-rejection.

Acknowledgments

This work was supported by grants NIH NIBIB R21 EB013716-01, NIH NCI R01 CA136569, NIH NCI R01 CA163498, and NIH NCI R41 CA180191.

The authors would like to acknowledge and thank Chengeng Zeng at the University of Washington’s Imaging Research Laboratory for helpful comments and discussion regarding the LQAF presented in this work.

References

Borghi G, Tabacchini V, Schaart DR. Physics in Medicine and Biology. 2016; 61(13):4904. URL: http://stacks.iop.org/0031-9155/61/i=13/a=4904. [PubMed: 27285955]

Espana S, Marcinkowski R, Keereman V, Vandenberghe S, Holen RV. Physics in Medicine and Biology. 2014; 59(13):3405. URL: http://stacks.iop.org/0031-9155/59/i=13/a=3405. [PubMed: 24888974]

Hunter WCJ, DeWitt D, Pierce LA, Morrocchi M, Miyaoka RM. Hardware implementation of principal-component-based signal multiplexing for a monolithic scintillator pet detector. 2017 manuscript in preparation.

Jorge L, Bonifacio D, DeWitt D, Miyaoka R. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment. 2016; 838:103–108. URL: http://www.sciencedirect.com/science/article/pii/S016890021630955X.

Ling T, Burnett TH, Lewellen TK, Miyaoka RS. PHYSICS IN MEDICINE AND BIOLOGY. 2008; 53(7):1843–1863. [PubMed: 18364542]

Ling T, Lee K, Miyaoka RS. IEEE TRANSACTIONS ON NUCLEAR SCIENCE. 2006; 53(5, 1): 2513–2518.

Ling T, Lewellen TK, Miyaoka RS. PHYSICS IN MEDICINE AND BIOLOGY. 2007; 52(8):2213–2228. [PubMed: 17404465]

Maas MC, Schaart DR, van der Laan DJJ, Bruyndonckx P, Lematre C, Beekman FJ, van Eijk CWE. Physics in Medicine and Biology. 2009; 54(7):1893. URL: http://stacks.iop.org/0031-9155/54/i=7/a=003. [PubMed: 19265203]

Marcinkowski R, Mollet P, Holen RV, Vandenberghe S. Physics in Medicine and Biology. 2016; 61(5): 2196. URL: http://stacks.iop.org/0031-9155/61/i=5/a=2196. [PubMed: 26907952]

Miyaoka RS, Li X, Hunter W, Pierce IILA, McDougald W, Kinahan PE, Lewellen TK. IEEE TRANSACTIONS ON NUCLEAR SCIENCE. 2011; 58(5, 1):2244–2249.

Miyaoka RS, Ling T, Lewellen TK. IEEE TRANSACTIONS ON NUCLEAR SCIENCE. 2007; 54(5, 1):1561–1567. [PubMed: 19180248]

Morozov A, Solovov V, Alves F, Domingos V, Martins R, Neves F, Chepel V. Physics in Medicine and Biology. 2015; 60(10):4169. URL: http://stacks.iop.org/0031-9155/60/i=10/a=4169. [PubMed: 25951792]

Pearson K. PHILOSOPHICAL MAGAZINE. 1901; 2(7–12):559–572.

Pedemonte S, Pierce L, VanLeemput K. Physics in Medicine and Biology. 2017; 62(21):8376. URL: http://stacks.iop.org/0031-9155/62/i=21/a=8376. [PubMed: 28436919]

Pierce LA, Hunter WCJ, Haynor DR, MacDonald LR, Kinahan PE, Miyaoka RS. Physics in Medicine and Biology. 2014; 59(18):5347. URL: http://stacks.iop.org/0031-9155/59/i=18/a=5347. [PubMed: 25146849]

Roweis ST, Saul LK. Science. 2000; 290(5500):2323–2326. URL: http://science.sciencemag.org/content/290/5500/2323. [PubMed: 11125150]

Schaart DR, van Dam HT, Seifert S, Vinke R, Dendooven P, Lohner H, Beekman FJ. 2008 IEEE Nuclear Science Symposium Conference Record. 2008:3581–3585.

Solovov VN, Belov VA, Akimov DY, Araujo HM, Barnes EJ, Burenkov AA, Chepel V, Currie A, DeViveiros L, Edwards B, Ghag C, Hollingsworth A, Horn M, Kalmus GE, Kobyakin AS, Kovalenko AG, Lebedenko VN, Lindote A, Lopes MI, Luscher R, Majewski P, Murphy ASJ, Neves F, Paling SM, da Cunha JP, Preece R, Quenby JJ, Reichhart L, Scovell PR, Silva C, Smith NJT, Smith PF, Stekhanov VN, Sumner TJ, Thorne C, Walker RJ. IEEE Transactions on Nuclear Science. 2012; 59(6):3286–3293.

van Dam HT, Seifert S, Vinke R, Dendooven P, Lohner H, Beekman FJ, Schaart DR. IEEE Transactions on Nuclear Science. 2011; 58(5):2139–2147.

Zernike, vF. Physica. 1934; 1:689–704.

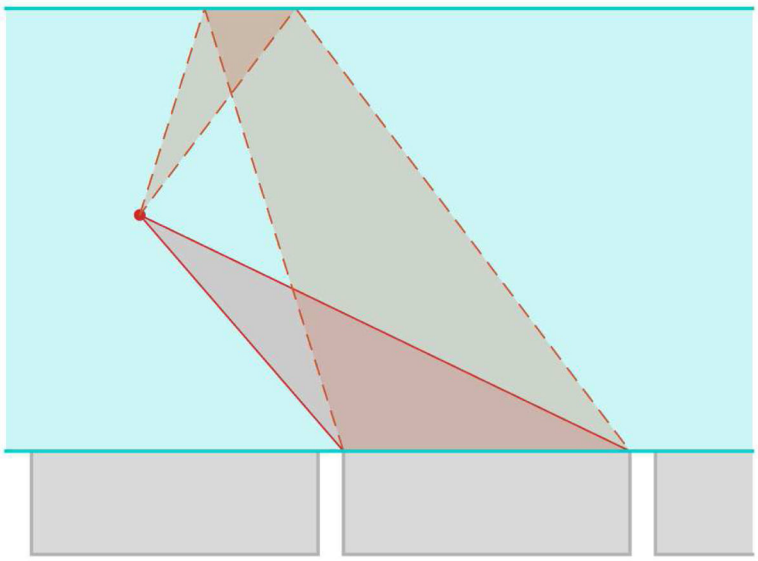

Figure 1.

An illustration of the simulated detector response setup. A photon interaction point is shown as the red dot, the scintillation crystal is shown in blue, the photosensors in gray. For a given photon, each detector output was modeled only as the solid angle from the photon interaction point to the detector (solid red cone), plus the reflection of the point to the detector (dotted red reflected cone).

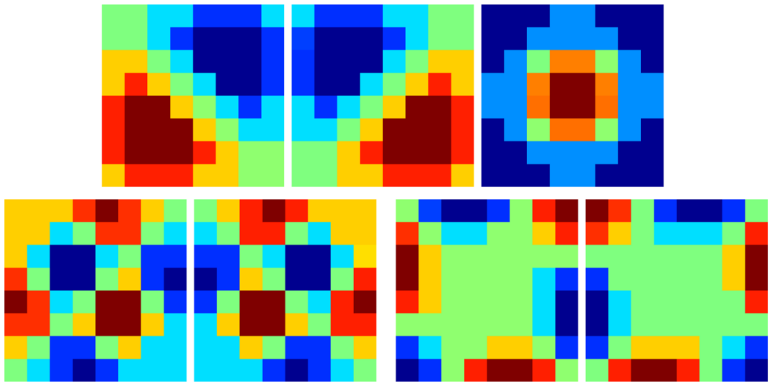

Figure 2.

The measured pulse weights from the multiplexing circuit. The top row are principal components 1, 2, 4; the bottom row shows the output from principal components 8 and 9 at left, with the two edge-specific channels at bottom-right. Note the similarity to the Zernicke polynomials (Zernike 1934).

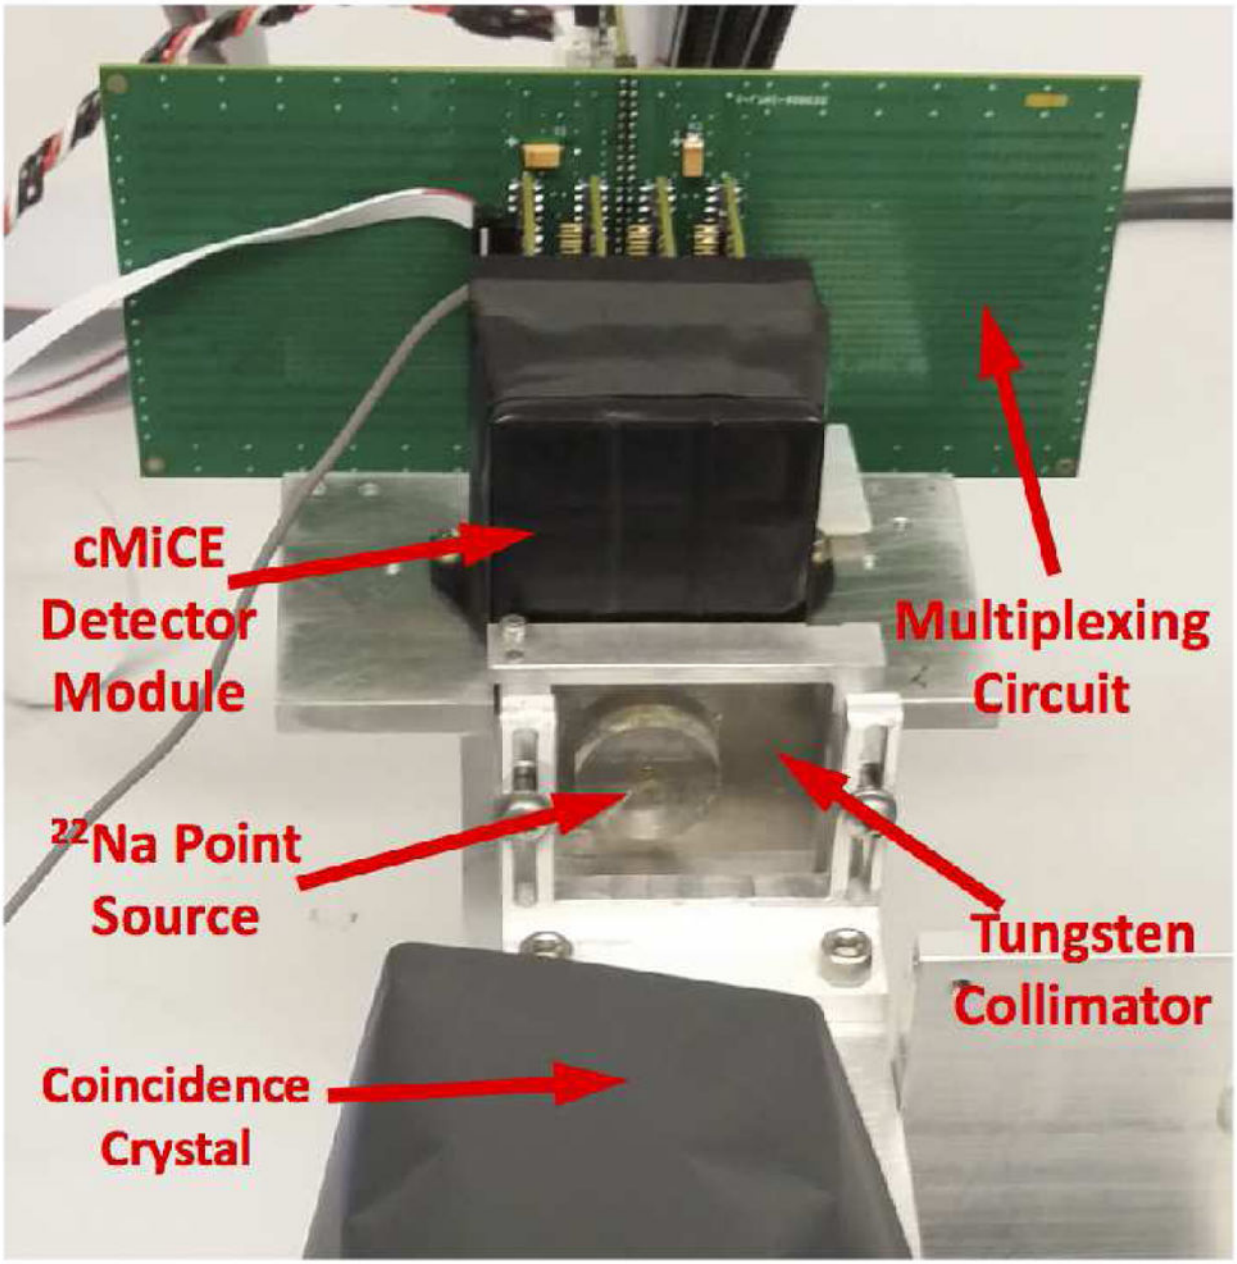

Figure 3.

Image of the calibration setup for the PCA-MUX multiplexed cMiCE detector module. During calibration, the ^{22}Na point source was 10 mm from the cMiCE crystal face, the coincidence crystal was 110 mm from the point source. In this image, the point source has been moved away from the detector face for illustrative purposes. A tungsten shield with a 2.0 mm hole allows the collimated beam to pass through, but prevents the cMiCE detector crystal from being flooded with random events.

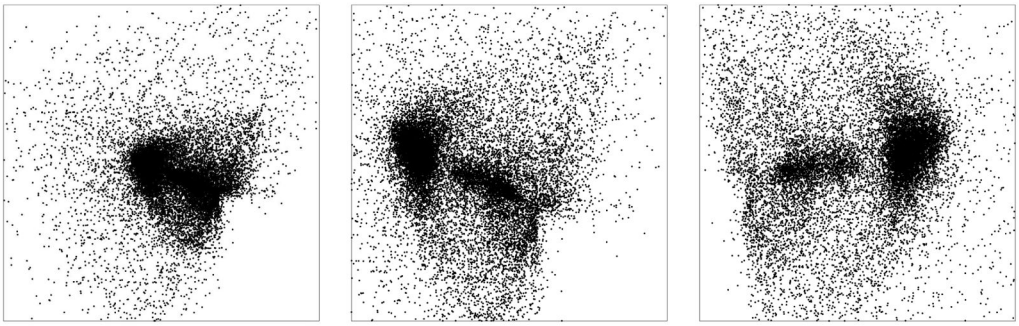

Figure 4.

(LEFT) A scatterplot of all 7-channel data events for a single calibration beam projected onto the first two most significant principal components. **(CENTER)** Same as the left image, but zoomed in. Two dense clusters are apparent amid the scattered events: a linear cluster (the “shaft of the arrow”) represents events that are low energy: down-scatters from the collimator or photons that down-scattered out of the crystal. The 511 keV photoelectric interactions lie within the “arrowhead” cluster. The Compton edge is apparent as the gap between the shaft and head of the arrow. **(RIGHT)** The same data after an initial round of scatter-filtering. The depth-direction is now more closely aligned with the x-axis (the vector that we call $PC_{2,1}$ in this section), with deeper-penetrating events to the right. The separation between the Compton edge and the photoelectric events in the arrowhead is now more pronounced. A projection onto the x-axis yields the $PC_{2,1}$ density distribution shown in Figure 5(b). The initial scatter rejection only serves to help align these two clusters and make the depth-direction the most significant principal component. This alignment was visually confirmed for all 1,089 beam positions.

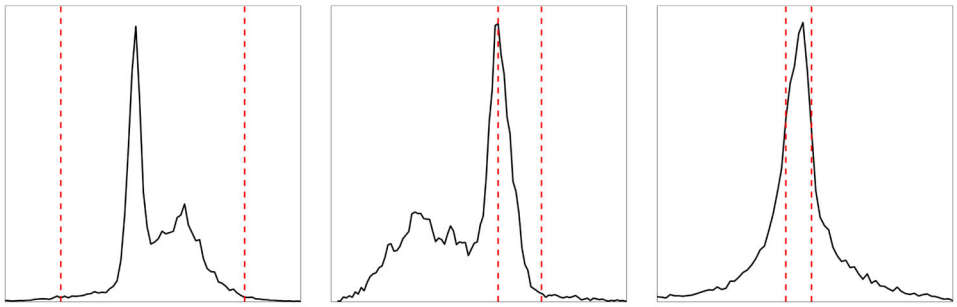

Figure 5.

(Left) The density profile along the first principal component of all calibration beam event data (the same data shown in Figure 4(a)). Thresholds (shown in red) were defined along each of the principal components as a fraction of the maximum density, events outside these bounds were considered scattered. **(Center)** The density profile along the primary principal component after initial scatter rejection (the same data as Figure 4(c)). This direction is used to estimate photon interaction depth (the “depth-direction”). This distribution is loosely correlated to the energy spectrum, but it is different. The events at the peak of the distribution are nearest the entrance surface and have lower energy. Events are considered deeper-penetrating as we move to the right. Note that the x-axis is non-linearly related to energy and depth, thus the distribution does not appear exponential to the right of the peak. The Compton edge is visible just to the left of the peak density. **(Right)** Density profile along a direction orthogonal to the depth-direction. Thresholds are defined as a fraction of the maximum density, events outside this threshold are rejected as scatter.

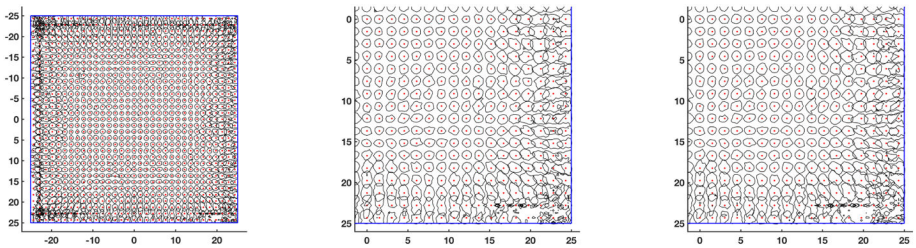

Figure 6.

(LEFT) Half-Max contours for each of the 33 × 33 test data beam positions for the 7-channel, 6-depth experiment. Red dots represent the nominal source positions. **(CENTER)** Same, but zoomed on the lower-right quadrant. **(RIGHT)** Lower-right quadrant half-max contours for the 7-channel, 10-depth experiment. The blue line represents the edge of the crystal. A region of bias is apparent as a blank space along the bottom edge near the right corner.

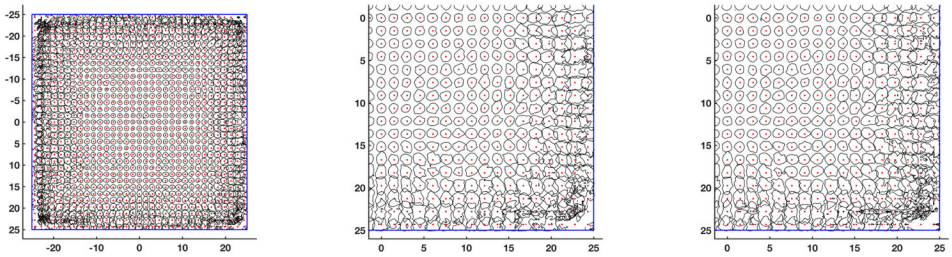

Figure 7.

(LEFT) Half-Max contours for each of the 33×33 test data beam positions for the 5-channel, 6-depth experiment. **(CENTER)** Same, but zoomed on the lower-right quadrant. **(RIGHT)** Lower-right quadrant half-max contours for the 5-channel, 10-depth experiment. The blue line represents the edge of the crystal. All four corners exhibit positioning bias for the 5-channel experiments.

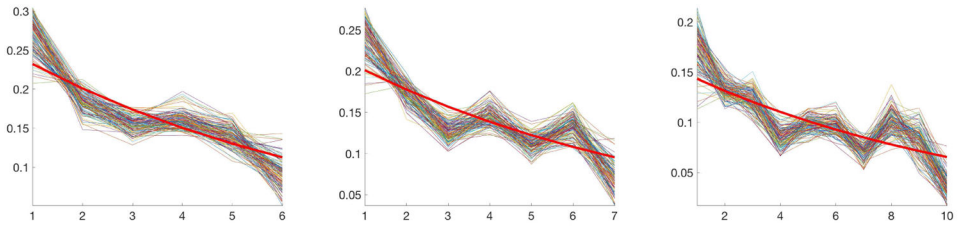

Figure 8.

Distribution of the number of events positioned per depth bin for the 7-channel experiments. **(LEFT)** For 6 depth bins. **(CENTER)** For 7 depth bins. **(RIGHT)** For 10 depth bins. Each horizontal axis indicates the bin depth number (with 1 being the entrance face of the crystal). The vertical axes indicate the fraction of events in a given depth bin. The distribution for each beam in rows/columns 12 through 22 (a square around the center of the crystal) are shown as the thin colored lines. The theoretical Beer-Lambert distribution is shown as the thick red line superimposed on each.

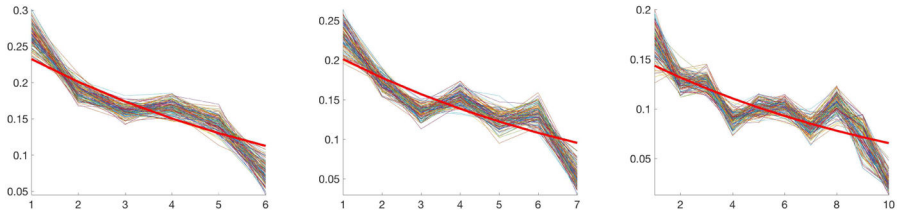

Figure 9.

Distribution of the number of events positioned per depth bin for the 7-channel experiments. **(LEFT)** For 6 depth bins. **(CENTER)** For 7 depth bins. **(RIGHT)** For 10 depth bins. Each horizontal axis indicates the depth bin number (with 1 being the entrance face of the crystal). The vertical axes indicate the fraction of events in a given depth bin. The distribution for each beam in rows/columns 12 through 22 (a square around the center of the crystal) are shown as the thin colored lines. The theoretical Beer-Lambert distribution is shown as the thick red line superimposed on each.

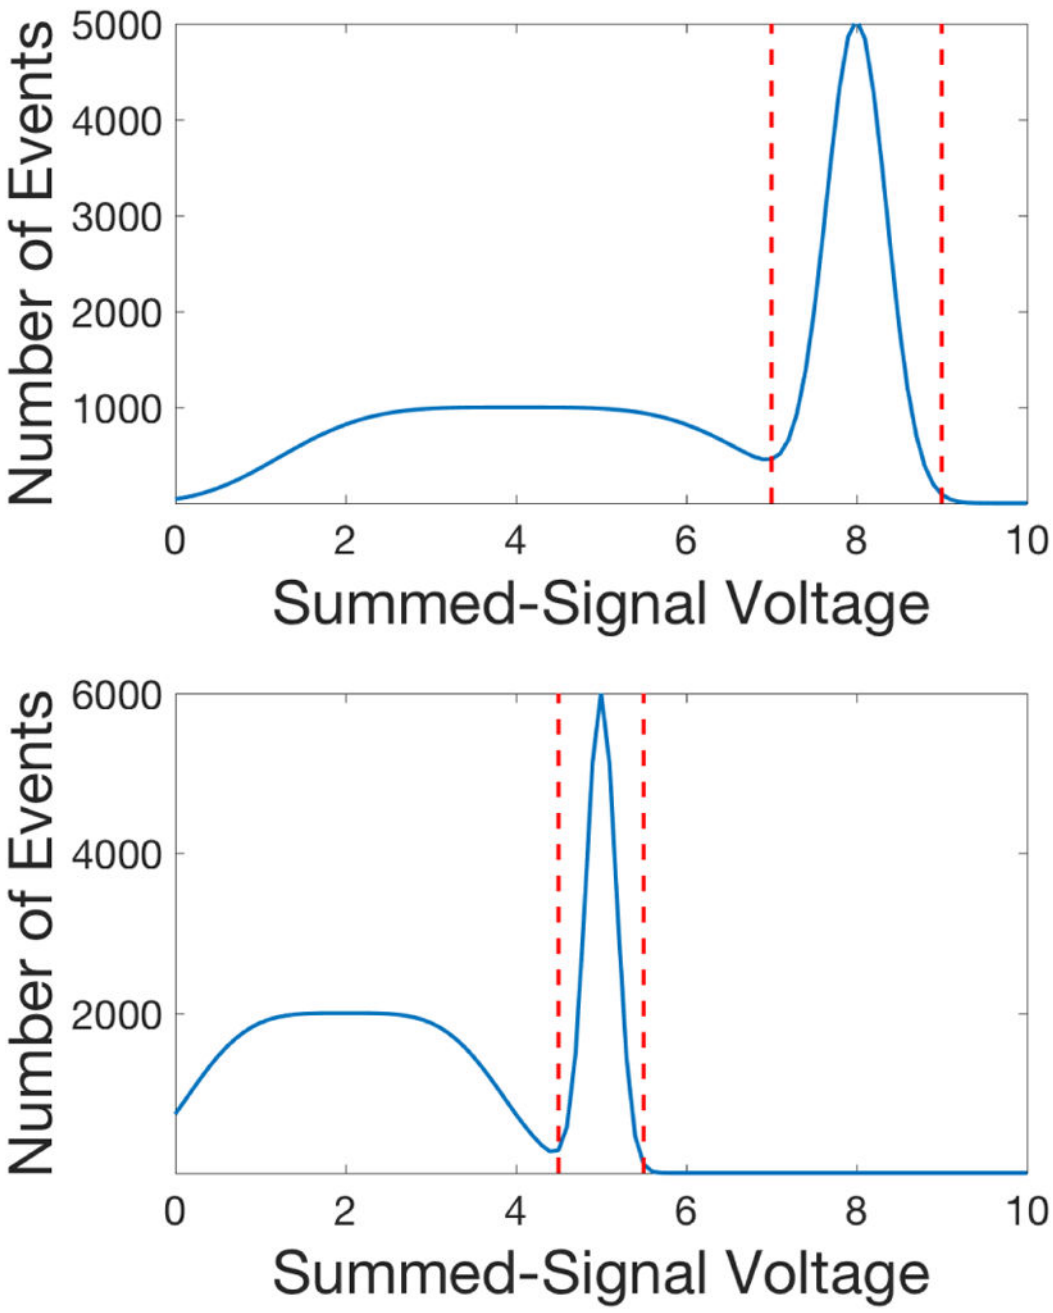

Figure 10.

Illustrations of summed-channel histograms from two different positions within a detector. The different positions of the photopeaks result in different energy windows that allow for low-energy scattered photons to be mis-positioned and accepted by the detector.

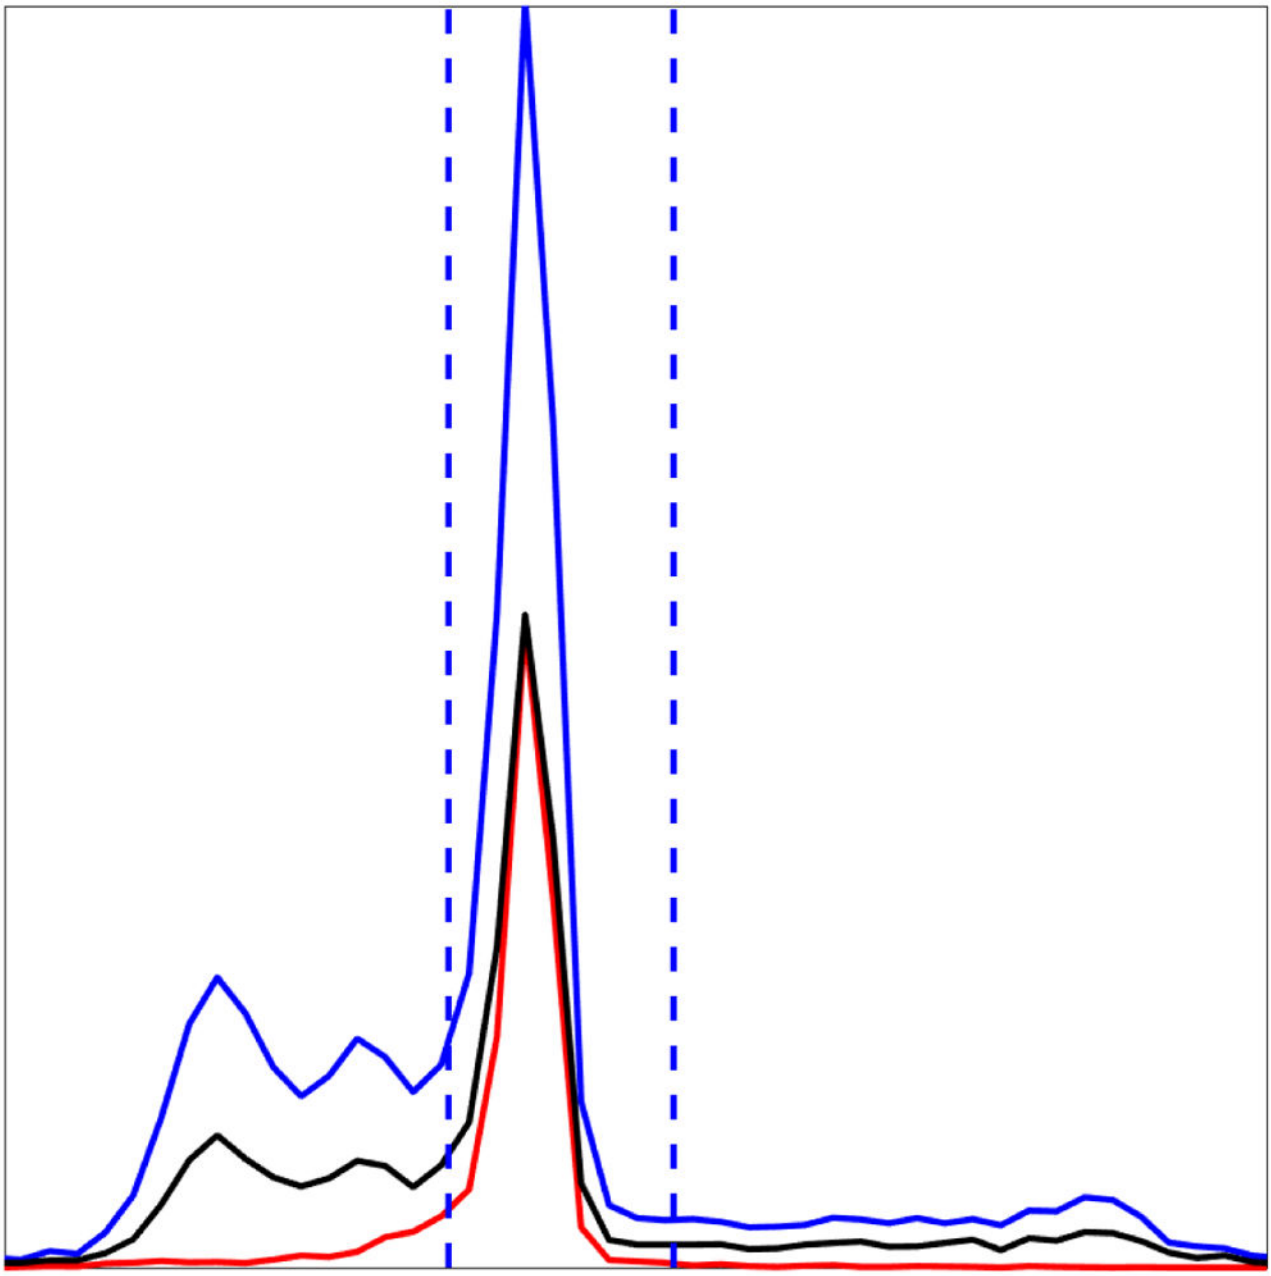

Figure 11.

Summed-channel energy histograms for a single beam position (row 13, column 17, near the center of the crystal) for the 7-channel, 6-depth experiment. The horizontal axis represents the voltage value output by the summed channel, the vertical axis represents the number of photon interactions with the given energy output. The blue line shows the energy distribution of all calibration data events from this beam position. An energy window (shown as dashed blue lines) was computed using this distribution. The testing data was positioned using the PC3 method and the events that were not scatter-rejected (those that were accepted as valid photon interactions) were used to create the energy histogram shown in red. The black line is the energy histogram of all testing data events that were accepted (not scatter-rejected) using

the initial-positioning scatter rejection method. The locally qualified acceptance fraction (LQAF) is the fraction of the black (or red) distribution that lies within the energy window shown (the correct energy window). The results for other beam positions and experiments are very similar.

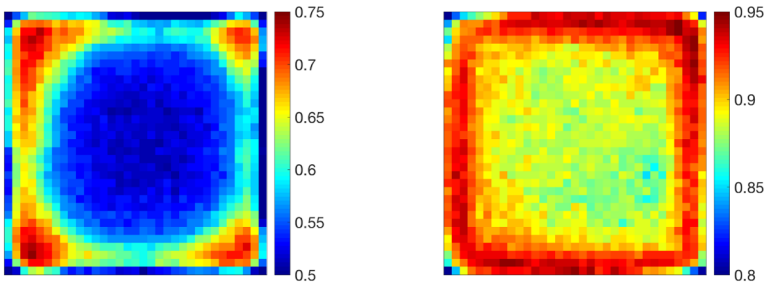

Figure 12.

Color maps of the locally qualified acceptance fraction for each beam position for the 7-channel, 6-depth experiment. **(LEFT)** Using the summed-channel method (from Figure 11: the fraction of the black distribution within the energy window defined by the blue distribution). **(RIGHT)** Using the proposed PC3 scatter-rejection method (From Figure 11: the fraction of the red distribution within the energy window defined by the blue distribution). Note the difference in color scales. Also note that the top-left beam position (row 1, column 1), was an outlier for both images, having a value of 0.40 for the left image and 0.74 for the right image. We increased the lower bound of each color map to better capture the dynamic range of the images. For the PC3 method, the scatter-rejection thresholds become increasingly strict (rejecting more photons) for the five outermost rows of beams.

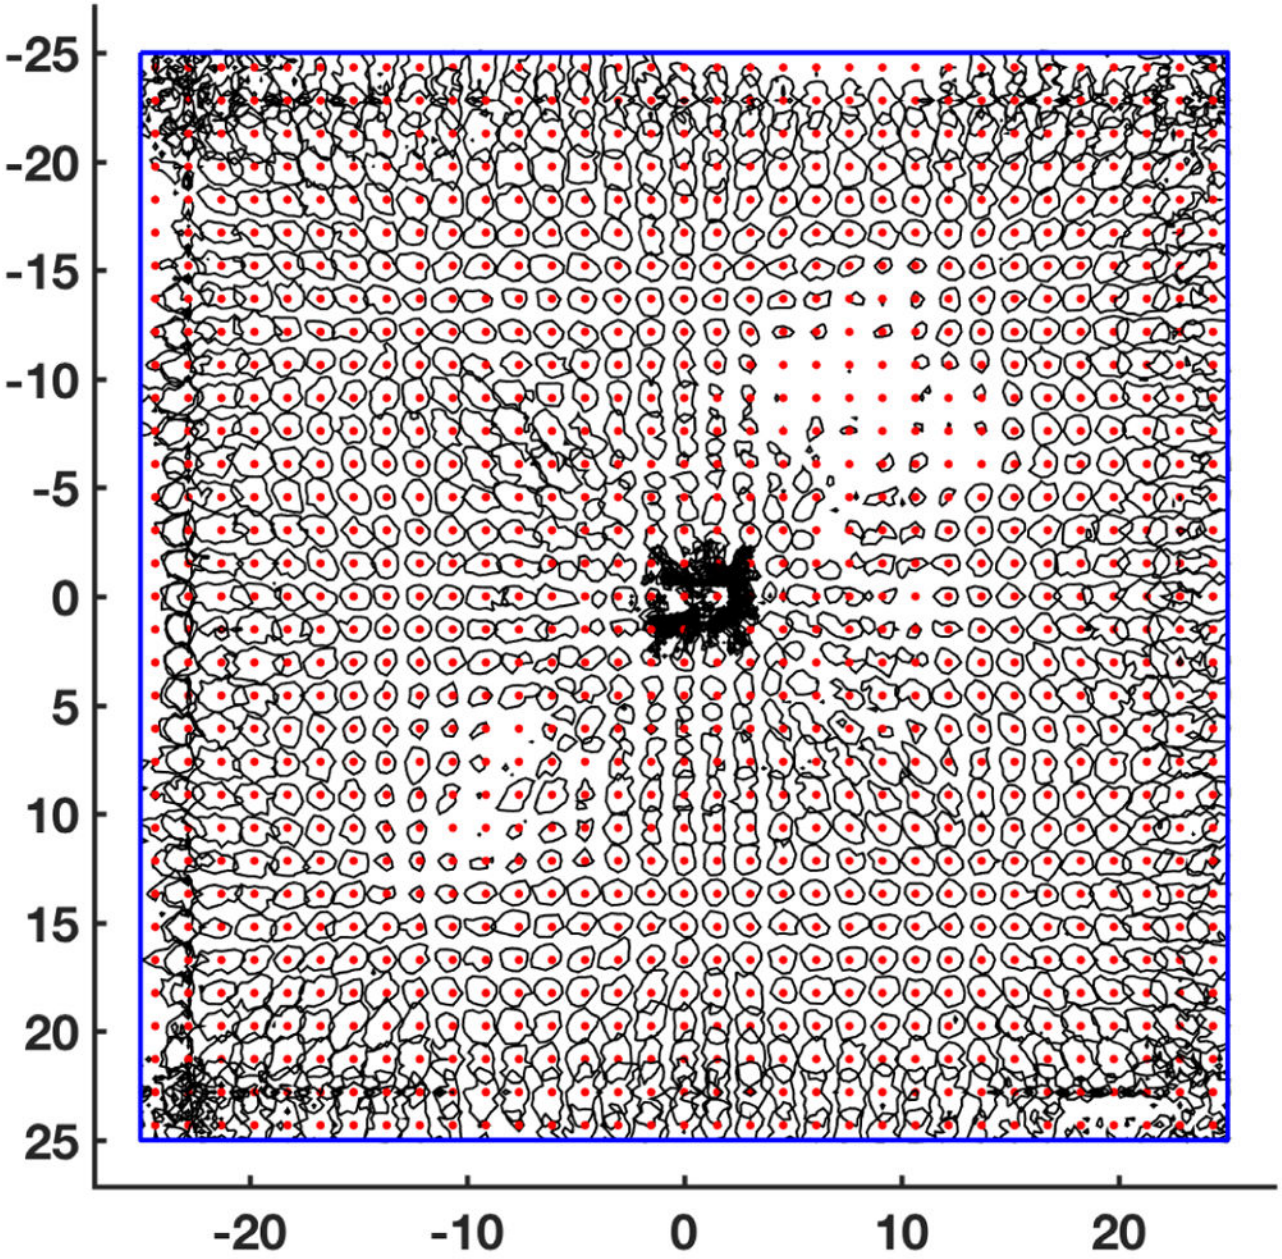

Figure 13.

Half-max contours for a poor choice of principle components. In this example, output channel 3 (principle component 4) has been excluded from the 7-channel, 6 depth experiment. This contour map was examined visually and this configuration of principle components was rejected by visual inspection alone, even though the quantitative FWHM measurements showed a mean FWHM of 1.3 mm.

Table 1

Results of the six positioning experiments (*ch* is the number of output channels and *D* is the number of depth bins for the experiment). These correspond to the half-max contour maps like those in Figures 6 and 7. Each measurement is given as the mean ± the standard deviation in millimeters. These results are from the contours only, and do not consider the 0.9 mm diameter beam width. Thus, these measurements are not the intrinsic resolution of the detector, which will be better than the numbers shown in this table.

		Center		Edge		Corner	
ch	***D***	**FWHM$_x$**	**FWHM$_y$**	**FWHM$_r$**	**FWHM$_t$**	**FWHM$_r$**	**FWHM$_t$**
7	6	1.17 ± 0.17	1.14 ± 0.17	1.89 ± 0.52	1.25 ± 0.25	1.73 ± 0.71	1.51 ± 0.59
7	7	1.18 ± 0.18	1.15 ± 0.17	1.89 ± 0.55	1.25 ± 0.27	1.93 ± 0.70	1.40 ± 0.50
7	10	1.18 ± 0.18	1.15 ± 0.17	1.88 ± 0.57	1.27 ± 0.25	1.98 ± 0.74	1.41 ± 0.62
5	6	1.23 ± 0.23	1.20 ± 0.19	1.94 ± 0.46	1.44 ± 0.29	1.71 ± 0.51	1.74 ± 0.59
5	7	1.25 ± 0.23	1.20 ± 0.20	1.97 ± 0.48	1.46 ± 0.30	1.66 ± 0.65	1.64 ± 0.65
5	10	1.25 ± 0.23	1.20 ± 0.19	1.99 ± 0.49	1.50 ± 0.32	1.71 ± 0.73	1.76 ± 0.62