Passive Gamma-Ray Imaging as a Verification Tool for Nuclear Weapons Treaties

K.P. Ziock, I. Garishvili, F. Gonzalez, K. Schmitt
Oak Ridge national Laboratory, Oak Ridge, TN[1]

Abstract

Treaties that include nuclear weapons as treaty accountable items (TAI) will invariably have to grapple with the question of how to verify that the TAI is indeed a nuclear weapon or a nuclear weapon component. While this may be a question that can be readily addressed with sufficient technical information about the TAI, the reality is that many of the technical details that allow an item to achieve nuclear yield are considered too sensitive for distribution, even among prospective treaty partners that are already nuclear weapons states. Nevertheless, to enforce such a treaty one will need technical means to inspect a TAI and verify that it is a nuclear weapon or weapon component. In particular, to provide confidence to the inspecting party, technical means must address TAI attributes that indicate nuclear yield is possible. One such attribute is the distribution of nuclear and overlying materials, and these can be determined using passive, coded-aperture, gamma-ray imaging. This paper addresses the fundamentals of gamma-ray imaging, and how it might be applied in TAI inspections.

Introduction

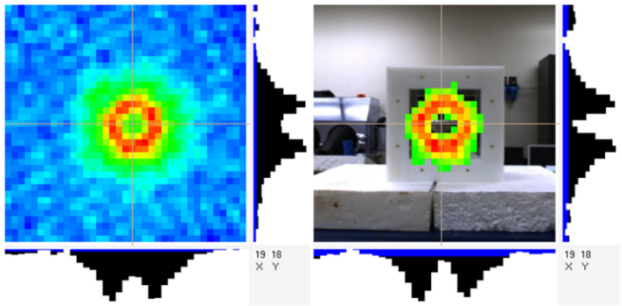

Fig. 1 Gamma-ray image of a depleted uranium storage casting inside a polyethylene box. The histograms to the right and below each image show the counts in the pixels under the tan cursor lines. The image on the right has the back panel of the box removed and shows much less scatter from the center of the cylindrical casting. The false-color threshold on the right image has also been raised to show the coaligned video image.

Coded-aperture, gamma-ray imaging provides a sensitive tool that allows one to form images of the fissile materials contained in nuclear weapons by imaging the passive gamma-rays emitted from those materials. In addition, because gamma-rays will scatter in nearby materials, losing some energy and being "re-emitted," such materials can also be imaged, and with an imager that has spectral resolution, the distributions of both radioactive and overlying materials can be obtained (Fig. 1). Such images provide a direct means to determine the distribution of such materials and therefore a direct measure of properties linked to the ability of the TAI to create nuclear yield. This capability makes gamma-ray imaging a candidate technical means for TAI authentication in future nuclear-arms-control regimes. However, with that capability comes the concern that information considered sensitive by TAI owners might be revealed. This raises the dilemma of the competing goals of the inspector and inspectee, with the former desiring a measurement with high certainty, while the latter will be concerned with the intrusiveness of an inspection. Of course, in a symmetric regime the roles of inspector and inspectee will regularly be reversed, so that any technical means used must be acceptable to all regime participants. Hence, before a given technique can be adopted, one must understand:

1) How the technique might be applied in an inspection,
2) The uncertainties associated with a given inspection,
3) How much sensitive information is at risk from an inspection (upon failure of any information barriers used to protect the information), and

4) How the data can be handled and analyzed so that sensitive information is protected, e.g. can the data be reliably processed behind an information barrier.

Integral to understanding these issues is an understanding of coded-aperture imaging and to that end a primer on this technique is provided below.

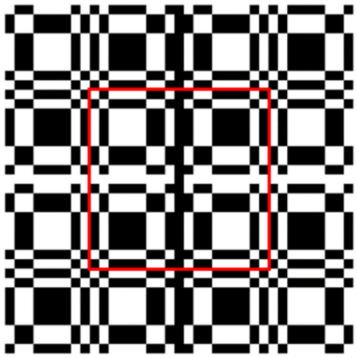

Fig. 2, Rank-19 coded aperture mask made of a ~ 4-fold repetition of the base pattern (inside the red square.) The region to the left of the square corresponds to the right side of the base pattern, the bottom to the upper part of the base pattern, etc.

Coded-Aperture Imaging

Coded-aperture imaging is an advanced form of pinhole imaging where the blocking sheet with the pinhole between the scene and the detector is pierced with multiple holes. This lets more radiation reach the detector, potentially increasing sensitivity, but comes at the expense of the images overlapping, which can both lower the sensitivity and create imaging artifacts if an arbitrary hole pattern is used. Patterns that overcome the latter problem have been known for several decades [1, 2] and these also allow calculating the statistical uncertainties in an image based on the counting statistics of the data [3]. The most common coded-aperture masks use a multi-fold repetition (typically 4-fold) of a base pattern that has a size comparable to the size of the detector (Fig. 1). When placed between a position-sensitive detector and the scene to be imaged, the sources in the scene cast the mask's shadow (shadowgram) onto the detector that encodes the scene as variations in counts across the face of the detector.

Modern mask designs are selected so that a point source in any pixel of the field of view casts a unique shadow pattern onto the detector, allowing one to determine the source's location [1]. They have the further property, akin to orthonormality, that the pattern projected by a second point source elsewhere in the field of view will not impact the ability to determine the location of the first source (and vice-versa). Mathematically, the image, $\boldsymbol{I}$, is recreated by cross-correlating ($*$) the mask pattern, $\boldsymbol{M}$, with the recorded shadowgram, $\boldsymbol{D}$:

$$\mathbf{I} = \mathbf{M} * \mathbf{D} \tag{1}$$

Or in terms of the individual matrix elements:

$$I_{i,j} = \sum_{m,n} M_{m-i,n-j} D_{m,n}, \tag{2}$$

where $I_{i,j}$ are the counts in image pixel *(i, j)*, $D_{m,n}$, are the counts in detector pixel *(m, n)*, and $M_{k,l}$ is the mask function at location *(k, l)*, defined to be 1 for open mask pixels and -1 for closed pixels. The mask provides good imaging properties based on the fact that any detector-sized portion (at location *m=i, n=j)* auto-correlates to a delta function, that is $I_{i,j}$ is 1 while all other locations are zero. Note that all counts in the detector are used to create each location of the image so that the statistical uncertainty can be shown to be equal to the square root of the total number of counts in the detector that are used to make the image [3]. That is a distinct difference from a direct imaging system where the image is formed directly on the detector so the uncertainty in each pixel of the image is only related to the number of counts in that pixel.

Regions of Interest

With a spectrally capable, position-sensitive detector, one can make an image using the data from each spectral bin of the detector to create a hyperspectral data cube; a structure that has an image linked to each spectral bin. From this structure, one can obtain the image associated with a spectral region of interest (sROI, e.g. a single or multiple emission lines) by adding the images from those energy bins (Fig. 3). One can also create the spectrum from a given area region of interest (aROI) by adding the

counts in the spectral bins associated with the included image pixels (Fig. 4). It should be noted that the spectral resolution of an imager is determined by the properties of its position-sensitive detector and are independent of the imaging. Instruments have been demonstrated with spectral resolutions ranging from none (from an integrating instrument based on scintillating glass fibers) [5] through the modest resolution of scintillator-based systems [6,7] to those with the excellent spectral resolution of high purity germanium [8].

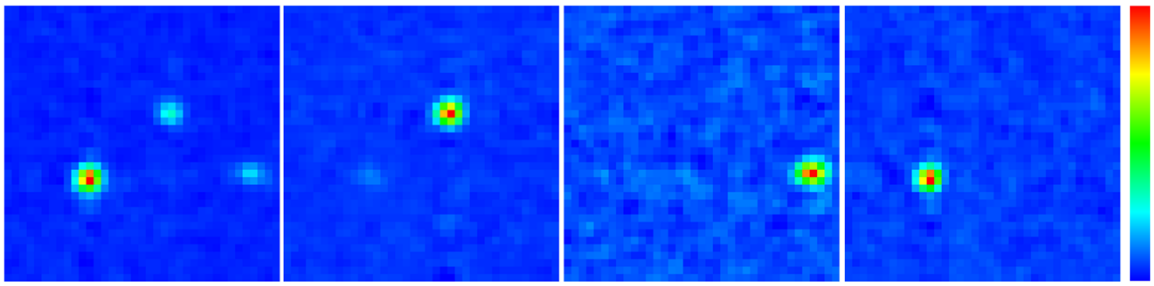

Fig. 3. Left, gamma-ray image of ^{133}Ba, ^{57}Co, and ^{232}U point sources with an overall sROI from 90-450 keV. The sources are individually located when specific sROI's are selected. From left-to-right, starting with the second image, sROI of 122-124 keV, 237-249 keV, and 354-358 keV locate the ^{57}Co, ^{232}U, and ^{133}Ba sources, respectively. The false-color scale applied individually from the minimum to the maximum of each image is shown on the right. The global spectrum is shown in Fig. 4.

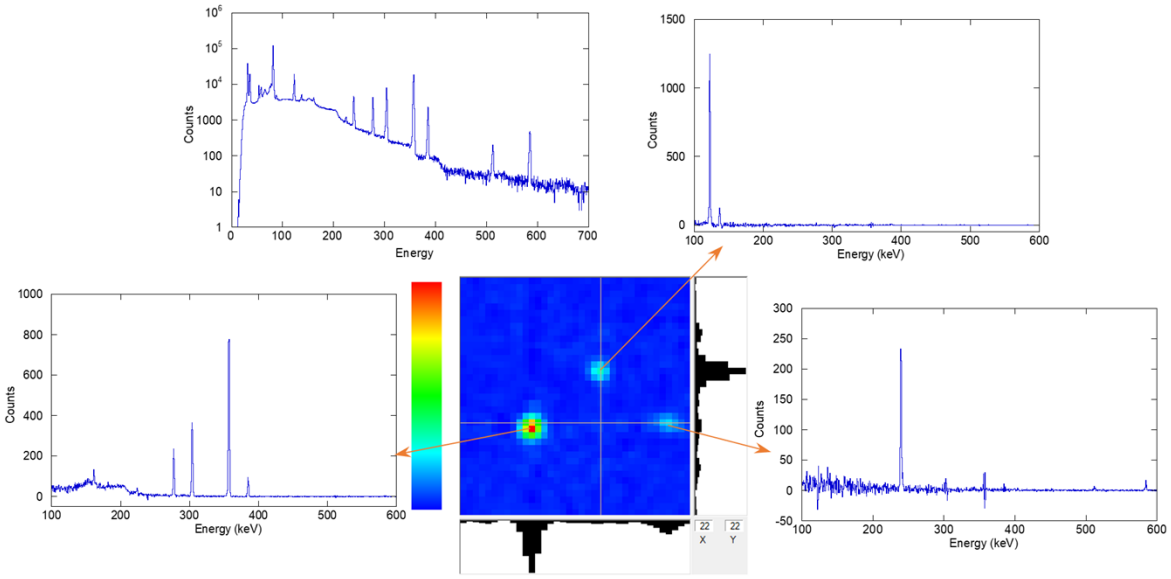

Fig. 4. Gamma-ray image of three point sources (bottom center) and the corresponding spectrum seen by the detector (top left). A single-pixel aROI at the center of each source in the image reveals the individual spectra for a ^{57}Co, ^{232}U (^{212}Pb daughter), and ^{133}Ba source (clockwise from top-right spectrum.)

Background subtraction

Because coded-aperture imaging works by encoding the scene as count-rate variations across the face of the position-sensitive detector, count-rate variations not due to the mask will create artifacts in an image. Such variations can be due to gradients in the radiation field from scatter off objects near the detector and they are always present. Sources of scatter include not only background radiation, but scatter from sources in the field of view itself. In addition, count rate variations can arise from residual detector non-linearities, and background sources outside the instrument's field of view. Fortunately, all such variations can be easily removed *in situ* by collecting data in two equal-time integrations, one with the mask, and one with its inverse, or anti-mask, where the open and closed holes are interchanged [9]. One can then add the two data sets to obtain the shape of the variations not due to the mask and subtract half

of that intensity from the mask data and half from the anti-mask data before generating each image. Those images can then be summed to form the final result:

$$\boldsymbol{I} = \left[\boldsymbol{D^M} - \frac{1}{2}(\boldsymbol{D^M} + \boldsymbol{D^A})\right] * \boldsymbol{M^M} + \left[\boldsymbol{D^A} - \frac{1}{2}(\boldsymbol{D^M} + \boldsymbol{D^A})\right] * \boldsymbol{M^A}. \quad (3)$$

Where the superscripts M and A refer to the mask and anti-mask, data and mask functions, respectively. However, because the mask function, $\boldsymbol{M^A}$, for the anti-mask is just the negative of the mask's mask function, $\boldsymbol{M^M}$, this can be simplified to:

$$\boldsymbol{I} = (\boldsymbol{D^M} - \boldsymbol{D^A}) * \boldsymbol{M^M}. \quad (4)$$

This technique makes coded-aperture imaging quite robust, providing the ability to obtain inherently background-subtracted data *without having to remove target sources from the measurement location*.

This is graphically demonstrated by the results shown in Fig. 4. The overall spectrum from the imager's detector (top) shows all of the peaks from the different sources. However, spectra obtained from aROI's containing each source individually, only have the peaks from that source. Close observation does show greater statistical variance at energies of the peaks from the other sources, agreeing with the previous statement that statistical uncertainty is proportional to the square root of all of the counts used to make a particular image.

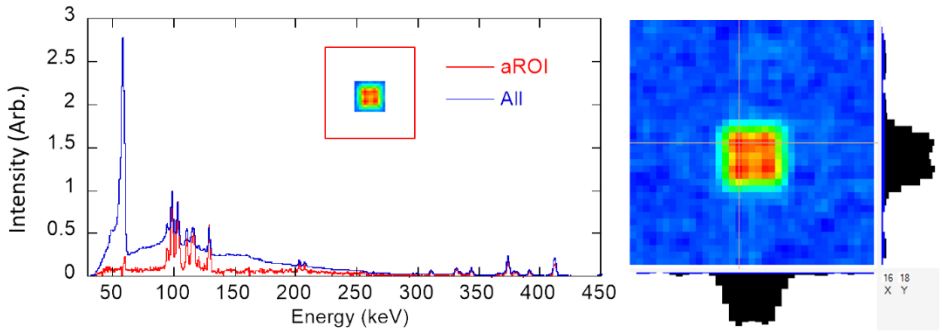

Fig. 5. Right, Gamma-ray image (2 hours of data) of 4 Pu Blocks arranged in a square configuration. Left, Spectra from the blocks as collected (blue) and with an aROI (inset, red.) The large 60-keV peak from the ^{241}Am daughter dominating the non-imaged spectrum is absent in the spectrum from the aROI because a Cu-Sn filter was used in front of the mask to block it. The latter spectrum also shows less down-scattered radiation below 250 keV. Note that the two spectra have each been normalized to their individual 98.5 keV K-shell fluorescence peaks.

The insensitivity to background is further demonstrated in Fig. 5 where a coded-aperture imager is used to image old ^{239}Pu blocks that have a significant daughter population of ^{241}Am. The image is obtained through a graded-Z, Sn-Cu filter designed to block all 60-keV radiation coming through the mask from reaching the detector. Never-the-less, the overall spectrum obtained by the imager's detector (blue) is dominated by the 60-keV peak from the ^{241}Am. However, if an aROI is selected to only allow data from the direction of the blocks (red), the 60-keV peak is removed. This shows that the observed 60-keV radiation reaches the detector by scattering around the mask, how normal background-subtracted data can be influenced by a source being measured, and the insensitivity of the coded-aperture technique to such background.

Image resolution

The spatial resolution, δx, of a-coded aperture imager is given by:

$$\delta x = \frac{a}{f} d \quad (5)$$

where *f*, the focal length, is the distance between the mask and the detector, *d*, is the distance from the mask to the object plane, and *a*, is base hole size in the mask. In essence, the resolution is limited to an angular change in the direction to the source, *a/f*, that shifts the pattern at the detector by one pixel of the mask. This is given as a linear dimension at the source plane by multiplying the angular resolution by

the distance to that source plane. Note that these are all properties that can be readily determined at the time of an inspection with a ruler.

However, it should be noted that this spatial resolution only determines the smallest structures that can be resolved. For instance, if a set of blocks is positioned in the field of view with each block much smaller than the resolution of the imager, then from the image alone, one could not tell if a single block had been removed. However, this does not mean that more detailed information could not be obtained from ancillary information. For instance, if one knows the emission rate per block, then with sufficient counting statistics, one could infer the number of blocks using a known count time, even without imaging. Similarly, if one has a model for the material distribution, for instance a square distribution given by the complementary error function in x and y:

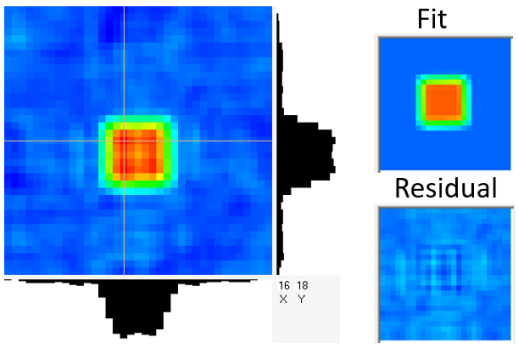

Fig. 6. Left, image of the 2 × 2 Pu block configuration (16 hours of data) from Fig. 5. Right, two-dimensional fit with expression 6 (top) and the residual (bottom). Some systematic noise is clearly present in the data.

$$N = B + N_0 erfc\left(\frac{x_R - W_x}{\sigma_x}\right) erfc\left(\frac{y_R - W_y}{\sigma_y}\right) \quad (6)$$

with:

$$x_R = |(x - x_0)cos\theta - (y - y_0)sin\theta|, \text{ and} \quad (7)$$

$$y_R = |(x - x_0)sin\theta + (y - y_0)cos\theta|,$$

then one could determine the overall width of the structure with a precision much better then δx by fitting to that model. This is shown in Fig. 6 where this square function is fit to the emission region of the 2 × 2 block array shown in Fig. 5. The data were collected in list mode, allowing one to resample the 16 hours of data on an event-by-event basis using boot-strap sampling (e.g. sampling with replacement) to generate many randomized trials with shorter integration times. The width of the fit in pixels as a function of time shows a statistical variation that decreases with increasing integration time and that is much less than the width of an individual pixel (Fig. 7), indicating that the statistical uncertainty on the fit width and height can be much better that the imager's resolution kernel.

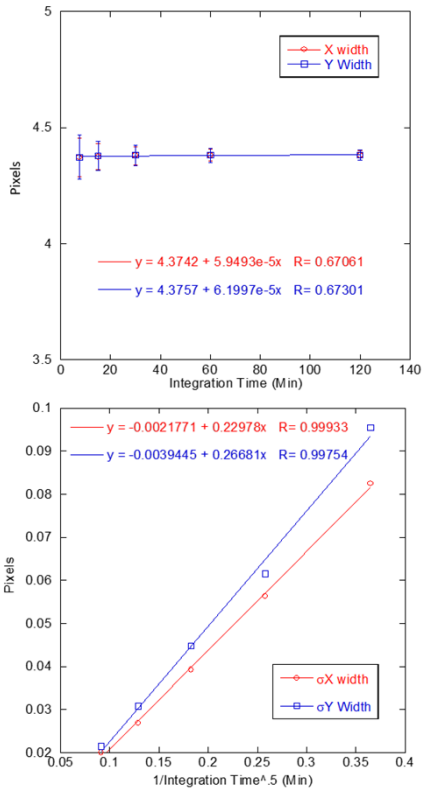

Fig. 7. Results of many boot-strap sampling fits to the 2 x 2 Pu block array as a function of integration time. Top, the width correctly gives the same value as the average of the overall data set, while the uncertainty is much less than a single pixel in size and decreases as the square-root of an increasing integration time (bottom).

Systematic biases

As a technique used to determine material distribution, one is interested in the accuracy of the results both for confidence in any given TAI inspection and to understand what sensitive data are at risk (e.g. details of material distribution) should an information barrier fail. This is an area that requires more research as indicated by the analysis of data collected on different block configurations with different imager resolutions. A sample of such data are given in Table 1 where the overall fits to the known number of blocks in each dimension are used to determine a width per block. As can be

seen, the dimensions generally agree for data collected with similar setups, but vary by more than the statistical uncertainties if the collection parameters change.

Table 1: Multi-block Fits

Config.	Block Size (Pixels)	Fit Size (Pixels)		Fit/True		Note
		X	Y	X	Y	
1 × 1	4.40	4.02	4.32	0.91 ± .002	0.98 ± .021	1
2 × 1	4.41	8.80	4.52	1.00 ± .006	1.03 ± .015	2
2 × 2	4.41	8.76	8.75	0.99 ± .004	0.99 ± .004	
2 × 2	4.44	8.73	8.76	0.98 ± .004	0.99 ± .006	
4 × 1	4.42	17.71	3.78	1.00 ± .004	0.86 ± .188	3
4 × 2	4.45	17.79	8.74	1.00 ± .002	0.98 ± .004	
4 × 3	3.48	14.70	11.14	1.06 ± .004	1.07 ± .004	4
2 × 2	5.75	10.41	12.48	0.91 ± .006	1.09 ± .008	5

Notes: 1) Single block is near the resolution limit, 2) 60 min of data, 3) Bimodal fit with ~20% having different y parameters, 4) Reduced Zoom, and 5) High zoom.

One obvious explanation for systematic biases is the property of skin depth. Unlike optical images where one is only imaging light from the surface of an object, the gamma-rays emitted by nuclear materials have a finite range in those materials. This means that the radiation intensity will change as one approaches the edge of the object which can result in a systematic bias that underdetermines the size of the object in a simple fit. It will depend on the energy of the radiation that is imaged, the attenuation length of that radiation, and the details of the surface geometry.

Imaged Spectral Fidelity

With a well calibrated imager, coded-aperture imaging can be quantitative as regards the flux emitted at the surface of an item. From images of thin (e.g. negligible self-attenuation), distributed U_2O_3 samples of highly enriched ^{235}U, mass estimates good to a few percent of the known ^{235}U mass were obtained from analysis of aROI-selected spectra. This was true for single items, the total mass from multiple items (single aROI), and individual items (each with its own aROI) in images with multiple items in the field of view [10]. This indicates that spectrally capable imagers could be used in place of simple spectrometers for some inspection scenarios. With the spatial resolution of the imager set so low that no TAI structure could be resolved, the imaged spectra provide the advantage that they are inherently background subtracted. With this approach, spectra could be collected with multiple objects in the field of view simultaneously, as well as in the presence of nearby nuclear materials that are out of the imager's field of view. However, we note that imaged spectra are inherently noisier (statistically) due to the imaging, and that sensitivity decreases at higher energies where the mask transparency increases. As with regular spectrometers, the spectra correctly measure surface flux and corrections for self-attenuation and attenuation by overlying materials must also be considered.

Information Barriers

Given the potential sensitivity of information collected by TAI inspection equipment, information barriers will likely be a necessary part of any system that is used in this capacity. It is an area of active research [11-14] and is not a unique requirement of imaging inspections so the interested reader is referred to the references for more information. However, there is one zero-information protocol that is unique to imaging that is worth mentioning because it does not inherently generate spatial information. Proposed by Marleau [15], the approach uses a template item that is otherwise certified as authentic and compares the imaging information on a unique event-by-event basis that does not integrate sensitive information. The template item is positioned on one side of the system while the comparable TAI under

inspection is placed on the other. The system itself comprises a central radiation detector surrounded by a cylindrical, rotating coded-aperture mask. The mask is designed with complete mask and anti-mask patterns such that the two are 180° out of phase. Consequently, at any given time, if the two items are identical, then the net exposure of the detector to the sum of the radiation from the two items is a constant. If the items are different, then the "mask" term will differ from the "anti-mask" term and a net temporal structure is seen in the count rate as the mask rotates. Thus, if no cheating occurs, then the system collects only total count rate information.

Potential Scenarios

Based on the preceding discussions, several broad classes of imager use for TAI inspections can be identified. Presented roughly in the order of sensitivity concerns (from the least to most sensitive) they are:

- Confirmation of the absence of SNM,
- Spectral TAI confirmation,
- Confirmation that a TAI has been altered, and
- Confirmation of a TAI.

The broad categories, and more specific hypothetical inspection scenarios suitable for coded-aperture imaging, are discussed below.

Confirmation of the absence of SNM:
Perhaps not an obvious class of applications, in these scenarios, a spectrally capable imager would be used to validate an inspectee's declaration that a particular item or region does not contain SNM. In this type of inspection there are no concerns about revealing classified information since there is not supposed to be a source of radiation to detect. In fact, the ability to image hidden SNM could serve as a strong deterrent to cheating, such as attempting to hide shielded SNM in a 'SNM-free' location. While this function could be performed with a non-imaging spectrometer, the imager provides several benefits. First, the imager's immunity to background from directions other than the field of view or objects being measured (imaged) increases the flexibility on where such work could be performed. In particular, it would not require performing such measurements at remote locations where no other SNM is present because flux from such material would not invalidate a null finding. Second, by using a well-characterized imager, one would be able to determine the minimum detectable activity for all pixels (or voxels) in the imager's field of view, irrespective of out-of-field SNM near the imager. Finally, the ability to generate images of illicitly located material and the concurrent revelation of sensitive information, could serve as an additional deterrent to discourage cheating.

Scenario 1, Non-nuclear confirmation: An item in an inspection region is declared as non-nuclear by the inspectee. A confirmatory measurement would be obtained by placing the item in front of the imager and data collected to validate that the item does not emit radiation (or at least radiation characteristic of SNM). Imaging advantages are:

- The imaging inspection could be conducted without a separate background-only measurement,
- The inspection could be conducted in regions where SNM is present without concerns about background lines contaminating the measurement,
- The uncertainty on the leakage flux could be determined from the counting statistics of the image data itself, and
- The imaging would serve as a deterrent to attempting to shield a hidden TAI. If down-scattered radiation from a shielded source in the item is detected, it could be imaged to determine the size of the shielding.

Scenario 2, Sweep source-free region: A given area could be declared source free, particularly an area where an inspection-related activity could or had occurred. Potential concerns could be hidden SNM that would be revealed in place of SNM that was supposed to be removed from a TAI, or hidden after removal. In this scenario an imager would be positioned to view the area and allowed to collect data for a time interval to validate the inspectee's claims that no source is present. Imaging advantages are:

- The imaging data are inherently background subtracted, something that could not be achieved with a traditional radiation detector since there is no item to remove to obtain a background measurement,
- The sequestered area could be located near SNM because the imaging data are inherently background subtracted, and
- The minimum detectable leakage flux at all locations throughout the area would be available, providing a quantitative determination of how much material could be hidden at each location.

Spectral TAI confirmation

Scenario 3, TAI spectral measurement: The absence or presence of SNM can be confirmed by measuring the gamma-ray spectra emitted by a TAI. This scenario is similar to Scenario 1, but in this case one seeks confirmation of the presence of SNM by observing a spectrum characteristic of such material. In this application, the use of imaging is based on its background rejection, rather than its spatial resolution. In fact, the measurement would be performed with the imager resolution set so low that structure within the TAI could not be determined. It should be noted that such measurements would deal with sensitive information and, like a measurement with a traditional spectrometer would require an agreement to share such information, the use of an information barrier, or both. The instrument's spectral resolution could be selected through the type of detector used in its construction. Imaging advantages are:

- The background-subtracted nature of imaging allows such measurements to be performed in proximity to other SNM, and
- The ability to obtain spectra from aROIs could allow multiple TAIs to be measured simultaneously. This requires that the TAIs have sufficient separation that the low spatial resolution of the imager required to avoid revealing individual TAI structure still allows non-overlapping aROIs.

Confirmation that a TAI has been altered

In these scenarios the goal is to confirm that the structure of a TAI has been altered as part of decommissioning or destruction processes. It encompasses a range of actions and confirmatory measurements targeted at both SNM and surrounding components [such as high explosive (HE)]. Scenarios could include observing the changes to the original TAI, or measurements on components that have been removed to validate that they either match parts of the original TAI, do not contain SNM, or both. It is envisioned that these measurements could resolve both spectral and structural detail raising sensitivity concerns that would require sharing sensitive information, the use of an information barrier, the use of a zero-knowledge approach, or a combination thereof. Confirmation of such measurements could be based on either first-principles or by comparison to measurements of a certified template item.

Scenario 4, TAI deconstruction, HE removal: The goal of this measurement would be to confirm that HE in close proximity to SNM was removed from the TAI. It relies on radiation from the SNM scattering in the HE, causing it to glow with gamma-ray emissions that allow it to be imaged. This measurement would involve placing the TAI in front of an imager configured with sufficient resolution to observe the spatial extent of the HE, should it be present. Ideally, the measurement would also have sufficient resolution to validate that the SNM present is of a distribution representative of the TAI prior to its disassembly. The advantage of imaging is:

- Simultaneous confirmation of both the absence of HE and the TAI-like presence of SNM.

Scenario 5, TAI deconstruction, SNM removal: The goal of this measurement would be to confirm that SNM has been removed from a TAI. It would ideally involve two measurements, one of the residual TAI confirming that it no longer contains SNM (see scenario 1), and a second to validate that the removed SNM matched a distribution commensurate with the original TAI. Such disassembly would likely occur in a region previously swept for the absence of SNM (see scenario 2). Ideally the measurements would have sufficient spatial resolution to confirm the SNM matched that expected for the TAI prior to its deconstruction (see scenario 7). Imaging advantages are:

- Multiple measurements could be performed with a single instrument (initial sweep of the measurement area for SNM, measurement of the residual TAI, and measurement of the SNM);
- The insensitivity to background means that confirmatory measurements of the lack of SNM could be performed without removal of the SNM from the immediate area;
- Confirmatory measurements of the lack of SNM could be performed to a known statistical precision; and
- The overall measurement location does not require a location otherwise free of SNM emissions. (e.g. it could be performed at one end of an SNM storage area).

Scenario 6, SNM destruction: Such a measurement would be used to confirm that the SNM distribution originating from a TAI had been sufficiently altered to preclude its reuse. In this scenario the TAI component would be measured with sufficient spatial resolution to confirm that its distribution had been altered to the point it could no longer function in its intended role. We note that sensitivity concerns of this measurement are lower than some of the preceding scenarios because the object should no longer meet the requirements of a nuclear weapons component. Imaging advantages are:

- Because material distribution can be related to nuclear yield, the measurement represents a direct probe as to whether or not a TAI was destroyed, and
- With concerns about revealing sensitive information, changes to the material distribution to the point that the component could no longer function, could be encouraged by performing a measurement with higher spatial resolution.

Scenario 7, TAI confirmation: in this scenario a TAI is measured to confirm that it is as claimed. A spatial resolution sufficient to resolve the spatial extent of components would be part of the verification. As an intrusive measurement, concerns about revealing sensitive information would need to be addressed. The advantage of imaging is:

- Direct confirmation of an attribute related to nuclear yield, from either first principles or via a template.

Conclusions

Because gamma-ray Imaging probes TAI characteristics that can be related to nuclear yield, its use as part of a technical inspection regime is being studied. As an indirect imaging technique, the first step in such a study requires understanding how it works and the potential strengths and weaknesses it entails. Primary among these are the ability to quantitatively image the distributions of source and overlying materials, spectral capabilities providing full hyperspectral imaging, and the use of mask-anti-mask imaging to provide inherent background subtraction. This latter capability means that the coded-aperture technique could be utilized to advantage in applications targeting spectroscopy. At the same time, there is a penalty associated with this form of imaging that means the signal-to-noise ratio from the image reconstruction needs to be considered in calculating sensitivities.

Based on the technique's unique capabilities, a number of TAI inspection scenarios have been identified where imaging could provide complementary benefits to other radiation detection techniques. These clearly require further refinement to tailor them to different classes and components of nuclear weapons. One requirement is to understand the limits of the sensitive information that could be revealed and that work is being undertaken using both data collection and Monte Carlo studies. Additionally, it is likely that any analysis will need to be performed behind an information barrier, limiting the ability of inspectors to validate each step of an analysis chain used to generate inspection results. These will need to convey minimal information, e.g. a pass/fail, while providing safeguards that the results are correct in an environment where inspectors cannot supervise each step of an analysis. A surprising outcome of this study are the scenarios where the imaging is not used to determine the spatial distribution of TAI components, but rather to confirm (quantitatively) the absence of SNM and the ability to obtain background-subtracted spectra in environments where other SNM and TAI's may be present.

References

1. E.E. Fenimore, T.M. Canon, "Coded aperture imaging with uniformly redundant arrays," Applied Optics, **17**, 337-347, 1978.
2. S.R. Gottesman, E.E. Fenimore, "New family of binary arrays for coded aperture imaging," Applied Optics, **28**, 4344-4352, 1989.
3. E.E. Fenimore, "Coded aperture imaging: predicted performance of uniformly redundant arrays," Applied Optics, **17**, 3562-3570, 1978.
4. M.C. Fleenor, M.A. Blackston, K.P. Ziock, "Correlated statistical uncertainties in coded-aperture imaging," Nucl. Inst. Meth. in Phys. Res., **A784**, 370-376, 2015.
5. GammaCam™ Radiation Imaging System, Innovative Technology Summary Report, OST Reference # 1840, (1998).
6. K.P. Ziock, et al., "A Gamma-Ray Imager for Arms Control Verification," IEEE Trans. of Nucl. Sci., **39**, 1046-1050, 1992.
7. K.P. Ziock, L. Nakae, "A Large-Area PSPMT Based Gamma-ray Imager with Edge Reclamation," IEEE Trans. Nucl. Sci., **49**, 1552-1559, 2002.
8. K.P. Ziock, C. Boehnen, J. Hayward, A.C. Raffo-Caiado, "A mechanically cooled, highly portable, HPGe-based, coded-aperture, gamma-ray imager*,"* Proc. of the 51st Annual INMM Conf., Baltimore, MD, July 2010.
9. J. Braga, et al., "A new mask-antimask coded-aperture telescope for hard x-ray astronomy," Experimental Astronomy, **2**, 101-113, 1991.
10. K.P. Ziock et al., Quantitative Holdup Determination Using Coded-Aperture Gamma-Ray Imaging," Proceedings of the INMM & ESARDA Joint Annual Meeting, Virtual, 2021.
11. S.M. Robinson, et al. "Imaging for dismantlement verification: Information management analysis algorithms, Nucl. Inst. Meth. in Phys. Res., **A662**, 81-89, 2012.
12. . M. Kütt, M Göttsche, A. Glaser, "Information barrier experimental: Toward a trusted and open-source computing platform for nuclear warhead verification," Measurement, **114**, 185-190, 2017.
13. F. Buhler, D.K. Wehe, M.P. Flynn, "A secure measurement unit for an inspection system used in nuclear arms-control verification," Nucl. Inst. Meth. in Phys. Res., **A982**, 164577, 2020.
14. C.J. MacGahan, et al., "Linear models to perform treaty verification tasks for enhanced information security," Nucl. Inst. Meth. in Phys. Res., **A844**, 147-157, 2017.
15. P. Marleau, R. Krentz-Wee, "Investigation into Practical Implementations of a Zero Knowledge Protocol," Sandia National Lab Report SAND2017-1649, DOI: 10.2172/1367490, available from www.osti.gov/biblio/1367490.