Imaging for Dismantlement Verification: Information Management and Analysis Algorithms

A Seifert   KD Jarman
EA Miller   AC Misner
MJ Myjak   WK Pitts
SM Robinson ML Woodring

September 2010

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

DISCLAIMER

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

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

**Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov**

**Available to the public from the National Technical Information Service,
U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161
ph: (800) 553-6847
fax: (703) 605-6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/ordering.htm**

 This document was printed on recycled paper.
(9/2003)

Imaging for Dismantlement Verification: Information Management and Analysis Algorithms

A Seifert   KD Jarman
EA Miller   A Misner
MJ Myjak   WK Pitts
SM Robinson   ML Woodring

September 2010

Pacific Northwest National Laboratory
Richland, Washington 99352

Imaging for dismantlement verification: information management and analysis algorithms

A. Seifert*, K.D. Jarman, E.A. Miller, A.C. Misner, M.J. Myjak, W.K. Pitts, S.M. Robinson, M.L. Woodring

Pacific Northwest National Laboratory, Richland, Washington 99352 USA

Abstract

The level of detail discernible in imaging techniques has generally excluded them from consideration as verification tools in inspection regimes. An image will almost certainly contain highly sensitive information, and storing a comparison image will almost certainly violate a cardinal principle of information barriers: that no sensitive information be stored in the system. To overcome this problem, some features of the image might be reduced to a few parameters suitable for definition as an attribute. However, this process must be performed with care. Computing the perimeter, area, and intensity of an object, for example, might reveal sensitive information relating to shape, size, and material composition. This paper presents three analysis algorithms that reduce full image information to non-sensitive feature information. Ultimately, the algorithms are intended to provide only a yes/no response verifying the presence of features in the image. We evaluate the algorithms on both their technical performance in image analysis, and their application with and without an explicitly constructed information barrier. The underlying images can be highly detailed, since they are dynamically generated behind the information barrier. We consider the use of active (conventional) radiography alone and in tandem with passive (auto) radiography.

Keywords: arms control, information barrier, treaty verification, warhead dismantlement

1. Introduction

The dismantlement verification process is defined by international treaties or other multi-party agreements. Inspectors might confirm that a small number of dismantled warheads are in a set of storage containers by making key measurements after these containers have been sealed. These measurements form the

---

*Corresponding author
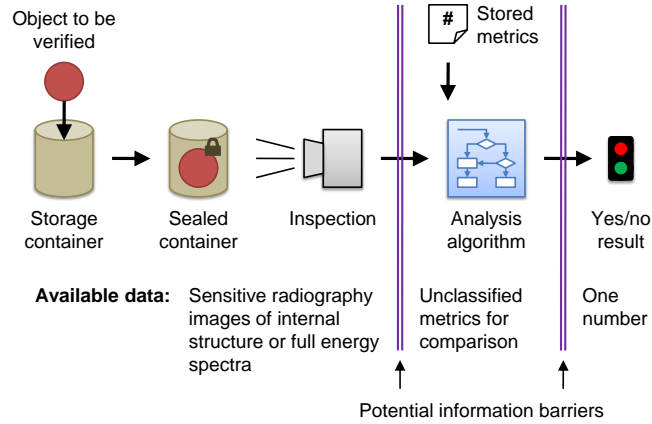*Email address:* allen.seifert@pnl.gov (A. Seifert)

Figure 1: Conceptual illustration of dismantlement verification process.

basis of comparison for later inspections of equivalent sealed containers. Preventing disclosure of sensitive information is critical for multi-party agreements. Hence, information barriers are typically established using physical obstructions, instrument firewalls, volatile data storage, rigorous procedural control, and other techniques. An information barrier (IB) is a collection of hardware, software, and procedures designed to prevent disclosure of sensitive information acquired during measurement.[1] The inspections must achieve a balance between protecting sensitive data and disclosing sufficient information to establish that the container holds the declared item. Figure 1 gives a conceptual illustration of the overall verification process.

Radiation imaging is not normally incorporated into the verification process. However, imaging gives the spatial relationship between subassemblies in the object to be verified, which provides one of the strongest diagnostics in assessing the intended function of that item. As the focus on Arms Control moves toward greater discrimination of items after the reduction in nuclear weapons from the Cold War era, the simple attributes defined for material control may not have sufficient fidelity to distinguish treaty-limited items, such as weapons, from non-limited items, such as components or stored material. At the same time, a dismantlement regime might utilize imaging to assess whether a subassembly is a weapon component, such as a pit, or just a lump of special nuclear material (SNM) having a similar radiation signature. Imaging might also have a role in a material control regime where SNM is stored in a particular declared configuration and/or chemical form.[2]

---

[1] As defined by the reports of the Authentication Task Force, chartered by the United States Departments of Defense and Energy, Authentication is the process by which a Monitoring Party gains the appropriate confidence that the information reported by a monitoring system accurately reflects the true state of the monitored item [1–4].

[2] A summary of the technical challenges facing Arms Control appears in [5].

Despite its advantages, imaging has been routinely dismissed as a verification technique, due in large part to assumptions that it is too invasive and too difficult to implement behind an IB. While imaging is invasive, its capability to assess form and function could be key to enabling future Arms Control agreements. Operating the analysis algorithm reliably and autonomously behind an IB presents a technical challenge; our work is gauging the magnitude of that challenge. Whether or not imaging is accepted or rejected is ultimately a policy decision, not a technical decision.

The main obstacle for using imaging is that simple comparison techniques require a stored reference image or set of parameters. There has been an implicit assumption that the stored image is sensitive. Storing sensitive information violates a cardinal principle of IB design. In addition, there has been a perception that other measurements, including high-resolution gamma-ray spectroscopy and neutron counting, give adequate information for many applications, such as SNM mass attribute measurements. However, as the total number of warheads is reduced far below historic highs, the time may come when assessing the function of a SNM item, rather than deciding whether or not the item is SNM in a given mass range, becomes the primary driver in a dismantlement regime.

In this paper, we first briefly summarize the two primary imaging techniques: active (conventional) radiography and passive (auto) radiography. Next, we discuss the utility of imaging in the context of Arms Control, and give some considerations for information barriers. We then present three potential image analysis algorithms that do not store nor disclose sensitive information about the item of interest. Finally, we give some initial results that illustrate the performance of the algorithms with various objects.

2. Imaging Techniques

A verification system can use radiography in both active and passive modes for scanning objects. In active radiography, a strong x-ray source is placed on one side of the object, and a large-area imaging detector is placed on the opposite side. An image is then made from the transmitted photon flux. The image formed on the detector is representative of the attenuation of the x-rays through the intervening materials (including the container and contents). Radiography techniques range from non-electronic dose-recording films to sophisticated imagers used for industrial applications. The shortcoming of active imaging is that it cannot determine the radioactive make-up of the constituent materials.

In passive imaging, or autoradiography, the radioactive object is used as the source for an imaging system such as a coded aperture imager. A number of techniques exist for this purpose [6, 7]; coded aperture imaging works well with the lower-energy gamma rays from SNM. In this technique, emissions from the target object are attenuated through a mask pattern and create a projection on a position-sensitive detector. The mask pattern is designed to render the detector response to a single point source as close to unique as possible [6, 8]. Coded aperture imaging has found use in imaging objects which are compact in the field of view, ranging in application from gamma-ray astronomy to national security

[9]. The drawback of passive imaging systems is that a solid mass of SNM could be replaced with a less dense object without significantly affecting the gamma-ray emissions outside the surface. In addition, the passive radiograph generally has significantly lower resolution than does an active radiograph. For instance, the resolution of coded aperture imaging is often limited by the size of the mask holes and the tradeoff between resolution and sensitivity.

Using both active and passive imaging systems may improve the likelihood of detecting material diversion, as it would be more difficult to "spoof" both systems at the same time. While passive radiography may be more effective at verifying the specifics of radioactive material distribution, active radiography may be more effective at verifying material density and structure.

3. The Utility of Imaging

The radiation image of a sealed container contains a large amount of information. The active radiograph can be processed to give the overall size, geometry, edge characteristics, and density of the constituent item. These parameters are interrelated. Edge characteristics and density should be related to each other as well as to the chemical form of the item, such as plutonium oxide or metallic plutonium. The geometry and edge characteristics should be related as well: for example, a plutonium cylinder (hockey puck) will have significant differences in edge attenuation between the curved sides and flat top or bottom. The passive autoradiograph generally indicates the distribution of radioactive material within the container. One can also extract position-dependent spectroscopic information from the passive imaging system.

Some examples of image information include the following:

- Geometry and shape: active and passive radiographs give an overall image of the item. Depending upon the quality of the images and characteristics of the imaging device, the information ranges from a simple high-contrast image showing the size and shape to a detailed image that can be analyzed for attenuation along the edges.

- Material composition: spectroscopic imaging with a gamma ray imager gives the overall isotopic composition of the item in addition to its size and shape. Passive radiography using a spectroscopic imager would give an image that could be analyzed to determine if a particular region of the image had the expected radiation signature. These techniques could also include a subset of active radiography using an external radiation source. This source can be used to induce nuclear reactions in the material, generating a secondary radiation signature [10, 11].

- Edge finding and density can be analyzed using conventional radiography with a high energy x-ray source, or with passive radiography using the inherent radioactivity of the SNM item. These parameters should be related back to both the size and shape of the item and to its mass.

- There is also a subset of imaging techniques in which spectra are compared along different sight lines to the object. This technique might give some of the same information in a less invasive manner than standard imaging. For example, if there are two or more SNM-containing regions, there will be self-occultation effects along particular sight lines. In a material control regime, for example, the declared form was two distinct but identical plutonium spheres stored one above the other in a common container. Comparing the measured count rates of the 375 keV and 424 keV plutonium gamma rays from the top and from the side would have given additional confirming information for the actual stored configuration.

4. Information Barrier Considerations

The other necessary component for an imaging system consists of the IB and associated operating procedures that protect sensitive data from disclosure during the inspection. The IB usually has several components: tamper-indicating enclosures around the measurement system to prevent unauthorized access, interlocks that shut down the system immediately upon tampering, and one-way data transfers that transmit the minimum required information. A key related concept is the reduction of measurements to a particular attribute with a defined yes/no outcome. For example, while a stored mass might be declared, the defined mass attribute might be a yes/no mass measurement within a related mass range. Reducing a sensitive measurement to an attribute simplifies the information barrier: the data output can be a simple yes/no result implemented with a one-way display. At the same time, information into the system is restricted by reduction to an attribute, thereby closing a possible path for tampering. The input data path might only be an identified path for a test or calibration item, for example. In all cases the flow of information is reduced to a minimum.

There is one important exception to the defined yes/no inputs and outputs through the IB. Hash functions are one-way transforms that take the entirety of a large message or file and reduce it to a condensed output message. It is fundamental that the hash function output cannot be used to recreate the original message. For authenticating the system memory and loaded software, a hash function can be implemented to assess the loaded software and memory by reducing an input string and total system memory storage to a defined hash output. The hash function is a component of the authentication process that assesses confidence in the overall system performance, including security considerations [12]. Implementing a hash function requires a method to input a data string and an associated output display system. Because the hash function is a one-way transform, there is no method to extract the sensitive state of the stored software and data on the system.

5. Analysis Algorithms

We explored several possible ways to deal with information security restrictions on the use of imaging in an Arms Control context, while maintaining the ability to distinguish between imaged objects. The first approach is to sufficiently obscure the image to the point where the information can be displayed. This process must remove the possibility of inversion in a manner analogous to hash functions. Another approach is to perform an internal analysis directly on the sensitive information behind the IB, returning a non-sensitive quantity; this alternative eliminates the need to store reference images or other sensitive parameters. A third approach is to compare active and passive radiographs of the same object behind the IB. This section presents three image analysis algorithms that were inspired by these respective information management strategies. We focused on two goals when developing the algorithms: the ability to distinguish between objects of different density, shape, and radioactive profiles, and the effectiveness of obscuring information more specific than the algorithms are designed to verify.

5.1. Histogram Comparison

The first algorithm works by computing a histogram of pixel intensities from an active radiography image. The histogram from the interrogated object is compared to a previously generated "template" histogram to determine a degree of similarity with the expected object. This histogram template would be stored on the system. We will address the question of information protection momentarily.

To study the capability for distinguishing between objects on the basis of pixel intensity histograms, we consider a set of image histograms with two or more groups of like images. We then estimate a comparison metric for every distinct pair of images, both for the within-group populations and the across-group populations. The average separation between within-group comparisons and across-group comparisons was considered to be the "resolving power" of this method, as greater separation suggests higher discriminating power. We used a two-sample statistical hypothesis test statistic to represent the average separation. Resolving power is quantified by an approximate statistical significance level.

We tested a number of histogram comparison functions. A simple dot product (evaluating the inner product between two vectors of histogram intensities) produced promising results and is the comparison function considered here. Let A_{ij} be an image for $i = 1, \ldots, M_j$ images in $j = 1, \ldots, N$ groups. Let X_{ij} be the vector representing the histogram of pixel intensities for image A_{ij}, and let $\rho_{ijkl} = X_{ij} \cdot X_{kl}$ be the dot product. Estimates of the within-group distribution mean and variance are as follows:

$$\overline{\rho}_{jj} = \binom{M_j}{2}^{-1} \sum_i \sum_{k>i} \rho_{ijkj}, \ s_{jj}^2 = \left[\binom{M_j}{2} - 1 \right]^{-1} \sum_i \sum_{k>i} \left[\rho_{ijkj} - \overline{\rho}_{jj} \right]^2.$$

This definition considers all unique pairs of non-identical images. Similarly, estimates of the across-group distribution mean and variance are as follows:

$$\overline{\rho}_{jl} = (M_j M_l)^{-1} \sum_i \sum_k \rho_{ijkl}, \ s^2_{jl} = (M_j M_l - 1)^{-1} \sum_i \sum_k \left[\rho_{ijkl} - \overline{\rho}_{jl}\right]^2.$$

We want to know whether across-group comparisons are significantly different from within-group comparisons. One way to answer this question is to use a two-sample t-test for significant difference between the means, accounting for the difference in variances [13]. The test statistic

$$\frac{\left|\overline{\rho}_{jl} - \overline{\rho}_{jj}\right|}{\sqrt{s^2_{jl} (M_j M_l)^{-1} + s^2_{jj} \left(\begin{smallmatrix} M_j \\ 2 \end{smallmatrix}\right)^{-1}}}$$

indicates the resolving power between groups of different images. Applying the t-test then provides a threshold against which this quantity can be compared to determine whether the resolving power is significant.

The question remains whether the histogram of pixel intensities is a sufficient one-way transform to fully protect sensitive information about the original image. A given image with $m \times n$ pixels can be rearranged in$(m \times n)!$ ways to produce other images with exactly the same histogram. This very large number of combinations makes any direct inversion prohibitive for all reasonable image sizes. However, it may be possible to extract some information about the imaged object from the intensity histogram. The fraction of "black/very dark" and "white/very light" pixels gives a measure of the overall size of the item relative to the imager's field of view. In addition, the fraction of "gray/transition" pixels give an overall estimate of the size of the edge. Using a lower number of pixels and a limited gray scale range limit the information that can be extracted. The limits on determining an object's size and edge transition can be found for a given bin size in the histogram. Once that technical result is known, the next step is determining whether the histogram binning scheme inherently protects sensitive information; if so, then the histogram should be viewable through an information barrier. We plan to pursue these questions in future work.

5.2. Material Recognition

Given an active image of an object, *a priori* knowledge of the object's general shape can be exploited to estimate attenuation characteristics and, ultimately, density of the material. The material recognition approach uses the pixel intensity gradients near the edges of the object to distinguish between lighter and denser materials. We developed algorithms for spherical objects as a relatively simple example. Figure 2 illustrates the overall concept behind the algorithm. The higher the material density, the sharper the edge transition.

In addition to knowing the shape of the object, a rough initial estimate of location and spatial dimensions is assumed. One way to obtain such an estimate is to use existing image analysis methods for locating and characterizing specific
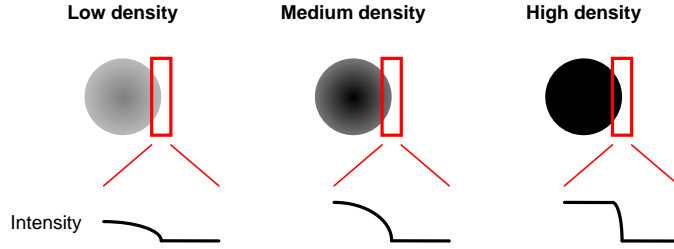
Figure 2: Determination of material density examining the intensity gradients at the edge.

shapes within images. The Hough Transform, for example, can be used to locate circles which could indicate spherical imaged objects [14]. Our initial testing of the circular Hough transform for this purpose showed promise, but is not the focus of the present study. Using Beer's Law and the parametrization of a sphere, we can then estimate attenuation factor by fitting that parametrization to a set of pixel intensities that correspond to some part of the sphere.

Let $I_{ij} = I_0 \exp\left(-\mu d_{ij}\right)$ be the intensity at pixel (i, j) within the sphere, where d_{ij} is the depth of the sphere at that pixel and μ is the attenuation factor for the material. The attenuation factor can be used as a discriminator of the type of material, given assumptions about the nominal densities of possible materials. An estimate for the unattenuated pixel intensity is assumed to be obtainable by a "blank" image or portion of the object image outside the sphere. In units of pixels, a sphere with center (i_c, j_c) and radius R is defined by

$$(i - i_c)^2 + (j - j_c)^2 + (d_{ij}/2)^2 = R^2.$$

Substituting for d_{ij} in the equation for the sphere and re-arranging gives the following equation:

$$[\log(I_{ij}/I_0)]^2 = 4\mu^2 \left[R^2 - (i - i_c)^2 - (j - j_c)^2\right].$$

Given pixel intensities and an estimate of no-attenuation pixel intensity, parameters (i_c, j_c), R, and μ are simultaneously estimated using nonlinear regression. The estimate of μ can be compared to empirical or simulated parameter values for a variety of materials to help verify the presence of a particular material.

Like the histogram comparison algorithm, the material recognition technique is designed not to disclose sensitive information. Although the radius of the object is computed as part of the process, only the material attenuation factor needs to be reported. The technique would be especially useful to distinguish between chemical forms since the metallic and oxide forms have very different densities. Hence, the output might be well suited to a metal/oxide attribute.

5.3. *Active/Passive Pixel Correlation*

The third algorithm studies the correlation between pixel intensities in the active and passive images of the same object in the same orientation. The key
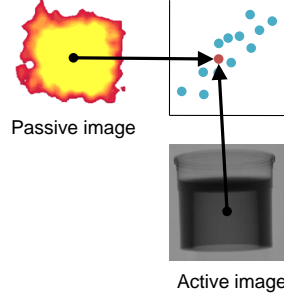
Passive image

Active image

Figure 3: Schematic for comparison of active and passive images for identifying dense, emissive material.

idea behind this algorithm is that the dense items in the container will be the most radioactive, whereas the surrounding filler material will be less dense and not emissive. (The passive image may contain some contributions from the filler material due to scattering.) For ease of comparison, we normalize both images to a set average pixel value, and re-bin the active image to match the number of pixels and corresponding spatial size of the passive image. (In general, active radiography systems are expected to have much higher resolution). The algorithm then groups each pixel from the active image with the corresponding pixel from the passive image, and records the pair of intensities in a scatter plot. Figure 3 depicts this process. The deviation of the scatter plot data from the diagonal indicates whether the object contains distributed, emissive material. If the dense items in the active image line up with the bright items in the passive image, the method should yield a strong correlation. If, however, the image contains less distributed materials or significant dense, non-emissive material (such as a point source along with a lead brick), the active and passive images will not match, and the correlation between image data will be much weaker.

Repeating the above description in mathematical terms, the algorithm transforms each active pixel value A_{ij} and passive pixel value P_{ij} so that the average of all the pixels in each transformed image is $1/2$. The value of x increases with higher density, whereas the value of y increases with higher emissivity.

$$x = -\frac{A_{ij} - \bar{A}}{A_{\max} - A_{\min}} + \frac{1}{2}, \; y = \frac{P_{ij} - \bar{P}}{P_{\max} - P_{\min}} + \frac{1}{2}$$

The algorithm then computes the slope of the linear regression model of y as a function of x. Again, if the images were identical in the sense of having a positive linear relationship between pixel values, this value would be 1, whereas values far from 1 suggest little similarity between the active and passive images. Other metrics could be used to evaluate the similarity, such as the average deviation of each scatter plot point from the diagonal.

The key advantage of the pixel correlation algorithm is the lack of stored parameters for image analysis. Hence, there are no information barrier restric-
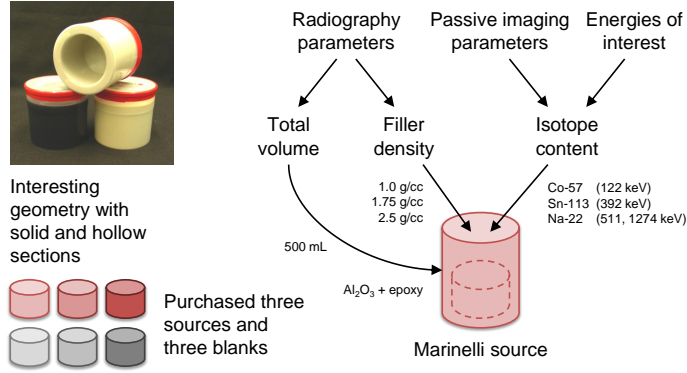
Figure 4: Marinelli beaker standards used for radiography.

Figure 5: Example radiographs of Marinelli beaker standards in different orientations.

tions to the technique. While initial results are promising, as presented later in this paper, further evaluation is necessary to consider a number of practical limitations: the strong attenuation of the x-ray flux in conventional radiography, the perhaps sparse gamma-ray flux available for autoradiography, and the possible presence of other high-Z materials, such as lead shielding in the case of stored plutonium oxide. One potential alternative is to have the inspection process use neutron flux to generate the active and/or passive images [10, 11].

6. Evaluation Methods

We evaluated the performance of the three image analysis algorithms on idealistic scenarios against a variety of real and simulated objects with different shapes, orientations, densities, and radioactive content. One set of objects consisted of six Marinelli beakers filled with epoxy having three different densities. One set of three beakers also included radioactive constituents homogeneously mixed in the epoxy. The total volume and densities were selected based on the parameters of the available radiography systems in the laboratory. Figure 4 gives the specifications of these objects. Orienting the Marinelli beakers in different directions produces interesting radiographs, as shown in Figure 5.

The beakers without included isotopes were imaged with a Varian PaxScan 2520 radiography system in various orientations. We used these images of

beakers without sources to evaluate the histogram comparison and pixel correlation algorithms. Simulated images of beakers containing radioactive isotopes were generated using MCNP [15] for a hypothetical passive imaging system. The model was based on the actual nuclides contained in the beaker as a source. The simulated passive imager consisted of an ideal pinhole camera with the aperture 100 cm from the front face of the beaker. The pinhole had a diameter of 0.5 cm and was 9.8 cm from the image plane. The image plane itself was 2×2 cm, containing 150×150 pixels. The resulting simulated active and passive images of the beakers were used for the pixel correlation algorithm.

A second set of objects was scanned with the active system to aid algorithm development, including rubber, quartz, and marble spheres of several sizes. We also created a third set of simulated objects for a hypothetical active imager based on the PaxScan. This dataset contained spheres of different materials and densities. We used both the actual and simulated sets of spheres to evaluate the material recognition algorithm. The simulated images were obtained utilizing a monoenergetic point source and a radiography tally in MCNP. The point source was placed 90 cm from the rear of the beaker and had an energy of 450 keV. The image plane was located 23 cm from the front face of the beaker. The dimensions of the plane were 30×30 cm^2 with 150×150 pixels. We included Poisson (counting) noise in all simulated images.

7. Results

7.1. Histogram Comparison

We tested the histogram comparison algorithm on a set of fifteen active images of Marinelli beakers: five images in three groups, with each group representing a different density of epoxy. The images in each group were taken under identical conditions, with a slight shift in spatial position of the beaker between images. Figure 6 depicts the fifteen images and corresponding histograms. These histograms were constructed by binning the pixel intensities into 5,000 equal bins. One can observe particular feature differences between the histograms in each group, as indicated by the red boxes in the figure.

Figure 7 compares the sample distributions of histogram comparison values between the low-density group and the low-, medium-, and high-density groups. The vertical axis shows the histogram comparison values generated by application of the dot product, in arbitrary units. A strong separation is indicated, and is supported by the results of the two-sample t-test described above. The significance level (in terms of a P-value) for a difference between within-group and across-group comparison values is below 10 to 15. A similar degree of separation was also found in the same kind of comparison on the medium- and high-density objects (not shown).

To check the overall methodology, we performed the same analysis with images randomly partitioned into three groups. The histogram comparison showed no similarity between these random groups, as expected. In more realistic scenarios, the level of variation in images of objects of interest is likely to be much
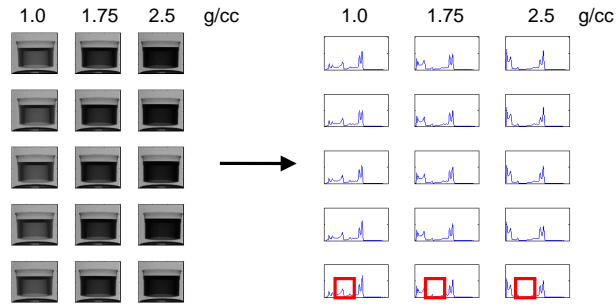
Figure 6: Marinelli beaker images and corresponding image histograms for evaluating the histogram comparison algorithm. The red boxes indicate a particular feature difference across object densities.
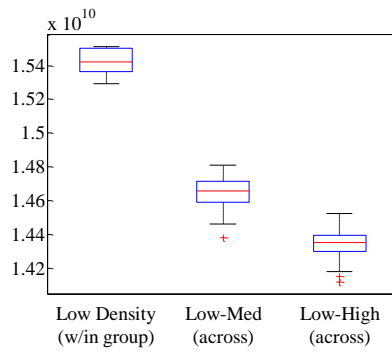

Figure 7: Box plot of within-group versus across-group histogram comparison values for the low-density Marinelli beaker images. Units are arbitrary. The across-group values are well separated from the within-group values.
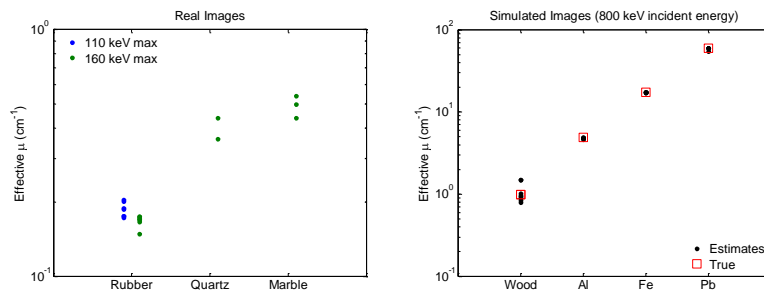
Figure 8: Results of the material recognition algorithm on real and simulated images of spheres. Each plotted point corresponds to a sphere of a different radius and material.

greater than that represented in the replicate measurements here, and further analysis on a broader range of image groups would be necessary to determine the full power of the approach. The results here show that the histogram comparison approach is very capable of distinguishing between the test objects of different densities and encourage such further study.

7.2. Material Recognition

We applied the material recognition algorithm to active images of real objects (including rubber, quartz, and marble spheres of several sizes) as well as simulated images of a range of materials. The active radiographs of real objects were taken with beam endpoint energies of 110 keV and 160 keV, whereas the simulated radiographs used a monoenergetic 800 keV or 1200 keV source. Figure 8 summarizes the results. Each point in the figure corresponds to a sphere of different radius and material.

The approach results in good discrimination between materials for both real and simulated images. The estimated μ for the large and small rubber spheres was very similar, while the estimated μ varied greatly between the rubber and marble spheres, even when these objects were similar in size. Even the densest objects considered in simulated images produced well-separated μ estimates for all seven sizes of spheres considered (noting the higher beam energy used for the simulations). Moreover, the estimates of this effective attenuation factor are very close to "true" values of the attenuation factors used for image simulation. This result supports further study of the material recognition algorithm for use in verification based on identification as well as discrimination.

Occlusion (where additional objects obscure regions of the one being examined) and regions of zero penetration (where the object is so thick or dense that the effective transmitted flux drops to zero) leave only portions of sphere images that can be used in this approach. We studied the effect of both these scenarios on the resulting μ estimates for simulated spheres of size 3 cm to 13 cm. As a baseline for determining minimal requirements, μ estimates were required to be within 10% of the calculated true value, and within 5% precision based on confidence intervals produced by the sphere fitting.
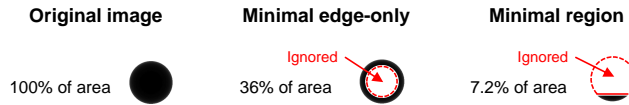
Figure 9: Minimum portions of simulated sphere images necessary for the material recognition algorithm.

Table 1: Correlation values for pixel correlation algorithm, relative to "end-on" active image of Marinelli beaker.

Passive image	Correlation
Beaker end-on	0.98
Beaker tilted 70 degrees	0.77
Beaker upright	0.67
Random field	0.01
Point source	0.65

Figure 9 depicts the minimum portions of the simulated images necessary for the material recognition algorithm to produce good results. The incident energy was 1200 keV in this case. Using only the outer ring (annulus) of each sphere image to estimate μ, the accuracy and precision requirements were met for the denser objects (Al, Fe, Pb) when the annulus consisted of only 36% of the full area for all sphere sizes. A lower percentage was needed for the larger spheres. We also studied the effects of occlusion by blocking out all pixels above a line through the sphere image. Again, the algorithm had good performance for denser objects at all sizes when only 7.2% of the full image could be used.

7.3. Active/Passive Pixel Correlation

To evaluate the pixel correlation algorithm, we took a real active image of the low-density Marinelli beaker in an "end-on" configuration. This image appears in Figure 5 at the far right. We then simulated passive images on the same beaker in the "end-on" configuration, tilted 70 degrees, and upright. Furthermore, we created simulated images of a random configuration (pixels uniformly distributed between 0 and 1) and a point source. Figure 10 illustrates the output of the pixel correlation algorithm. The plot on the left contains the noisy data from the images, whereas the plot on the right depicts the linear regression lines for the scaled data. The correlation values for these images are reported in Table 1. As expected, a near-perfect correlation is obtained from the active-passive comparison of end-on images, while a steadily decreasing correlation is found as the passive image varies, ultimately yielding no correlation for a random image.

The results quantitatively demonstrate the concept of pixel correlation and suggest that this approach may be useful as a verification concept.
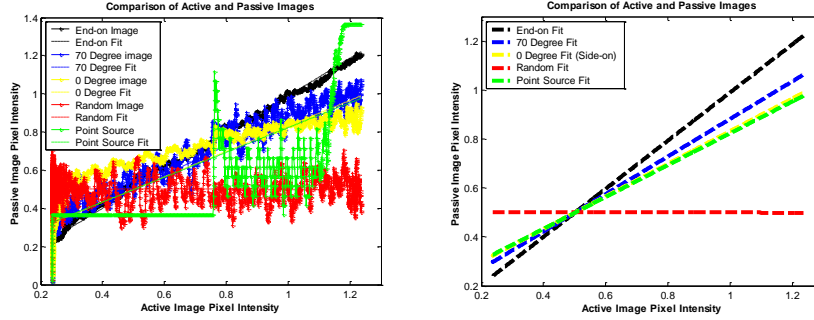
Figure 10: Results of pixel correlation algorithm. An active image of the Marinelli beaker in "end-on" configuration was compared to various simulated passive images. The plot on the left shows the individual correlation values, whereas the plot on the right gives the linear regression lines.

8. Conclusions

We have presented three possible techniques for analyzing imaging information in an Arms Control context. These algorithms are designed to avoid the tradeoffs between unambiguous object verification and applicability to environments where full images may not be retained. Each technique has well-posed and clear-cut technical questions to be addressed from an information security standpoint. Specifically, the histogram comparison algorithm has a question that can be rigorously addressed using mathematical analysis. The material recognition algorithm might be well suited to a metal/oxide attribute. The pixel correlation algorithm rigorously satisfies IB constraints since it does not rely on any stored information at all. We do not resort to explicit "blurring" or removal of image data to provide information security. Together with knowledge of the radiography systems, the use of these techniques, alone or in combination, can potentially improve verification capability and increase the likelihood of detecting material diversion.

The algorithms can be improved in several ways. For subtly differing images with a high degree of internal complexity, histogram comparisons seem to perform well. However, this may not be the case when imaged objects are simpler, or detector settings fluctuate or are unknown. Therefore, further study must consider these metrics against such images, and should develop other metrics to examine differences between simpler objects. The material recognition algorithm as it was tested presupposed the location of a circular object within the image and an accurate measurement of its location and radius. Additional noise is expected when this method is applied to images with objects of unknown size and location, and further work should provide a complete analysis of the resulting uncertainties. The pixel correlation technique has been performed on a variety of images, but its actual performance will depend strongly on the variety and sensitivity of the passive gamma ray imaging system used. A real passive system for inspection will need to employ an imaging system having resolution

at least similar to the expected number of spatial bins needed for the analysis methods outlined in this work. Furthermore, this technique presupposes minimal additional material or structure. Further research must determine its sensitivity when additional material is present.

Acknowledgment

References

[1] The Joint United States DOE-DoD Information Barrier Working Group, The functional requirements and design basis for information barriers, Tech. Rep. PNNL-13285, Pacific Northwest National Laboratory, Richland, Washington, Official Use Only (1999).

[2] The Joint DOE-DoD Authentication Task Force, Guidelines for authenticating monitoring systems, Tech. rep., Official Use Only (June 2001).

[3] R. T. Kouzes, J. L. Fuller, Authentication of monitoring systems for nonproliferation and arms control, in: Proc. Symposium on International Safeguards: Verification and Nuclear Material Security, International Atomic Energy Agency, Vienna, Austria, 2001, IAEA-SM-367/17/05.

[4] R. Whiteson, D. W. MacArthur, Information barriers in the trilateral initiative: Conceptual description, Tech. Rep. LAUR-98-2137, Los Alamos National Laboratory, Los Alamos, New Mexico (1998).

[5] M. R. Smith, J. Dunn, K. Seager, Future directions for arms control verification technologies, in: Proc. INMM 2010 Annual Meeting.

[6] E. E. Fenimore, T. M. Cannon, Coded aperture imaging with uniformly redundant arrays, Applied Optics 17 (3) (1978) 337–347.

[7] J. F. Morgan, G. Ignatyev, D. Semenov, M. Chernov, Gamma-ray camera for arms control applications, in: Proc. SPIE, Vol. 3769, 1999, pp. 24–30.

[8] A. Busboom, H. Elders-Boll, H. D. Schotten, Uniformly redundant arrays, Experimental Astronomy 8 (1998) 97–123.

[9] M. J. Myjak, J. S. Rohrer, S. J. Morris, M. L. Woodring, J. H. Ely, Pulse processing system for the RADMAP radiation modulation aperture imager, in: Proc. 2007 IEEE Nuclear Science Symposium.

[10] P. A. Hausladen, P. R. Bingham, J. S. Neal, J. A. Mullens, J. T. Mihalczo, Portable fast-neutron radiography with the nuclear materials identification system for fissile material transfers, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 261 (1–2) (2007) 387–390.

[11] J. T. Mihalczo, J. K. Mattingly, J. S. Neal, J. A. Mullens, NMIS plus gamma spectroscopy for attributes of HEU, PU and HE detection, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 213 (2004) 378–384.

[12] R. Hansen, B. Bass, R. Kouzes, N. Mileson, Implementation of the AES as a hash function for confirming the identity of software on a computer system, Tech. Rep. PNNL-14170, Pacific Northwest National Laboratory, Richland, Washington, Official Use Only (Jan 2003).

[13] P. J. Bickel, K. A. Doksum, Mathematical Statistics, Prentice Hall, Englewood Cliffs, New Jersey, 1977.

[14] R. O. Duda, P. E. Hart, Use of the Hough transformation to detect lines and curves in pictures, Communincations of the ACM 15 (1) (1972) 11–15.

[15] The X-5 Monte Carlo Team, MCNP—a general n-particle transport code, version 5, Tech. Rep. LA-UR-03-1987, Los Alamos National Laboratory, Los Alamos, New Mexico (1987).