AUTOMATED TEMPLATE MATCHING METHOD FOR NMIS AT THE Y-12 NATIONAL SECURITY COMPLEX

J. A. Mullens, J. K. Mattingly, L. G. Chiang, R. B. Oberer, J. T. Mihalczo

ABSTRACT

This paper describes a template matching method used by the Nuclear Materials Identification System (NMIS) to identify weapons components. The method is applied to NMIS's active source (^{252}Cf) radiation measurements for HEU components, utilizing four scintillation detectors. NMIS measures the spatial and temporal distribution of neutron and gamma radiation after a ^{252}Cf fission. This method further processes these measurements to extract pattern recognition features for the matching algorithm, and finds the closest matching component in the NMIS library of reference component features (templates). This identification method is being developed for use at the Y-12 National Security Complex. The goal of this development is to obtain high accuracy with the constraints of short measurement times and a small number of reference template measurements.

INTRODUCTION

The Nuclear Materials Identification System (NMIS) uses a ^{252}Cf source and a stack of two to four 4" plastic scintillation detectors to perform identification of HEU components in containers. The system uses a Windows PC to do data acquisition and analysis. NMIS works well for HEU components because it interrogates the component with high-energy neutrons and gammas from the ^{252}Cf source.

A template matching method matches the measurement of an unknown item to *templates* of known types of items. A template representing a type of HEU component is constructed by measuring several known components. Then during inventory confirmations, unknown components are matched to the template to confirm their type. Matching requires some analysis because of the natural variation of the measurement result from item to item. It allows shorter measurement times and some flexibility in equipment operation. Template matching is the logical method for inventory confirmations when templates can be created easily and stored securely.

This method's development draws on the growing collection of templates and confirmation measurements at Y12. The method is totally automated in the sense that the software loads the individual measurements of the template sets and automatically decides what data from the measurements should be used to make the matching decisions. This selection is done quickly enough that it is simply redone each time the program is started rather than stored for later use.

DATA AND SIGNATURES

The ^{252}Cf source is placed on one side of the HEU component container and the detectors are stacked on the opposite side. The NMIS data acquisition board measures the times at which each ^{252}Cf fission occurs and the times counts occur in each detector, with 1 nanosecond accuracy, for millions of counts per second. These signals describe the spatial and temporal distribution of neutron and gamma radiation coincident with the ^{252}Cf fission.

The ^{252}Cf and detector pulse data can be represented by a time series for each signal i, $x_i(t)$, that has the value 1 if a pulse occurred during the 1 nanosecond (ns) sampling interval and 0 otherwise. NMIS processes this data to produce the *correlation* between two signals i and j, $C_{ij}(\tau)$, for each block of T sampling intervals acquired:

$$C_{ij}(\tau) = \frac{1}{T-\tau} \sum_{t=0}^{T-1-\tau} x_i(t) x_j(t+\tau)$$

where t = time (sampling interval), and τ = delay of interest between detector i count and detector j count. $C_{ij}(\tau)$ is averaged over all data blocks collected. This correlation is actually a coincidence measurement, which NMIS measures simultaneously for all time skews τ between detector signals i and j. For example, if the time of flight of a gamma ray from the ^{252}Cf source to a detector is 3 ns, $C(3)$ represents the coincidence between the ^{252}Cf fission and the resulting gammas reaching the detector. (For a more detailed description of the correlations see [1].) Data is collected for several minutes, during which time roughly 100 million ^{252}Cf fissions might take place. Each fission is essentially a small experiment testing how the HEU component interacts with neutron and gamma radiation.

$C_{ij}(\tau)$ is a *signature* of the HEU component measured because the neutrons and gammas emitted by the ^{252}Cf source interact with the component as they pass from the source to the detector. NMIS calculates and uses other signatures, most of which have been used at some time for template matching. This application uses only the correlation between the ^{252}Cf source fission and each of the four scintillation detectors. Since these signatures require coincidence with the system's ^{252}Cf source fission, they remove other nearby sources of radiation.

The plastic scintillation detectors are sensitive to both gamma rays and neutrons. Matching uses the average $C_{ij}(\tau)$, which can be thought of as the rate of arrival of radiation at time τ ns after a ^{252}Cf fission. The first arrivals are any gamma rays that pass directly through the component at the speed of light; these are soon followed by scattered gammas. Directly transmitted high-energy neutrons arrive next, followed by lower-energy neutrons, scattered neutrons, products of induced fission in the HEU component, and any radiation scattered from the surroundings. Fig. 1 shows measurements on four aluminum blocks used to test the matching method. This measurement shows the arrival of gamma rays (1 ns) and neutrons (8 – 20+ ns).

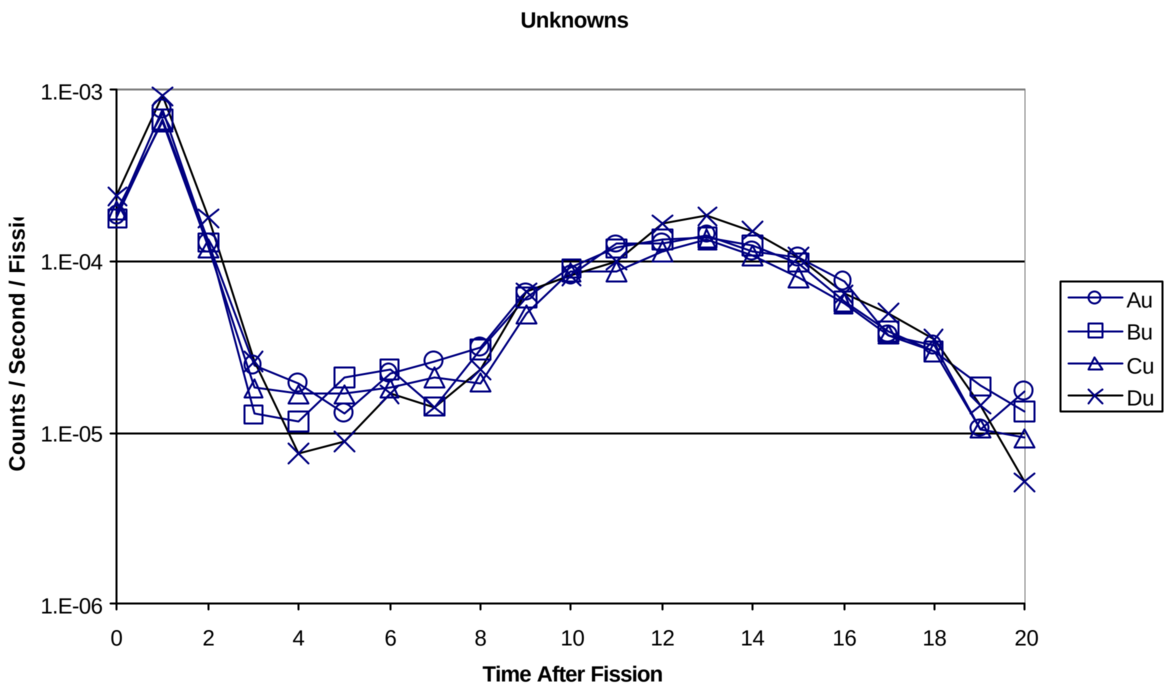

Figure 1 Correlations from four aluminum block measurements.

TEMPLATES

In this application, a template consists of a set of correlation signatures for some type of HEU component. The template matching algorithm identifies an unknown HEU component by comparing the unknown's correlation signature to all of the available templates, and selecting the closest matching template as the identity of the unknown. (The unknown has a declared identity, but the matching algorithm does not consider this.) The template set is treated as samples from a probability distribution of possible measurements of that type of HEU component. In reality the measurement has both random and deterministic elements. The major contributors to variations in the template set are:

- The randomness of the ^{252}Cf fissions and the radiation interaction with the component produce a natural variation. This is reduced by measuring for longer times, but not eliminated.
- The variations in manual positioning of the ^{252}Cf source and detectors relative to the component introduce more variation. After careful positioning of the NMIS detectors, there is still some variation of the component's position in the container.
- Measurement equipment calibration setup and drifts introduce some additional variation. These are effects are small since equipment calibration is checked daily.

The ideal template would contain statistically independent samples. This would require going through the entire process of setup, calibration, detector positioning, and measurement, using a different HEU item of that type and different operators, for each measurement in the template set.

A more expedient process was used to create the initial templates now in use, with some care to create these variations in the template set. It was decided to use only twelve fairly short measurements to construct a template (and several templates still have fewer). Statistically speaking, this is a small number from which to estimate the probability distribution of the measurements. The results have been satisfactory, but some templates have subsets of related measurements that clearly form a cluster. This sometimes affects statistical analyses of the template set.

For this application, a different template was created for every pairing of HEU component type with a container type. Also, measurements made in very different surroundings have separate templates since some ^{252}Cf radiation interacts with the surroundings then reaches the detectors.

MATCHING METHODS

The matching methods employed in previous efforts started with a simple Mahalanobis distance [2] comparison. The distance between each template signature and the unknown is in units of standard deviation of the template signatures:

$$d_{ij}^{2}(\boldsymbol{t}) = \frac{\left(C_{ij}(\boldsymbol{t}) - m_{ij}(\boldsymbol{t})\right)^{2}}{\boldsymbol{s}_{ij}^{2}(\boldsymbol{t})}$$

where i and j are the signal indices and $\mu_{ij}(\tau)$ and $\sigma_{ij}(\tau)$ are the template's mean and standard deviation of $C_{ij}(\tau)$. The total distance is the sum of d^2 over all τ and signal combinations ij selected by the analyst as signatures. The unknown is identified as the nearest template (smallest distance). This technique works reliably if μ and σ are known accurately and the unknowns' measurement have the same probability distribution as the template. But in practice, this technique has problems unless the estimate of σ is very accurate. It is especially sensitive to underestimating $\sigma_{ij}(\tau)$. This version of the Mahalanobis distance assumes that the $C_{ij}(\tau)$ are not correlated with each other, so that their covariance matrix is diagonal with each element being $\sigma_{ij}(\tau)$. The full covariance matrix version naturally has even more problems. [shorten this?] The software employs several techniques to make the estimate of $\sigma_{ij}(\tau)$ conservative, if selected by the analyst:

- Set σ to the value required by the randomness of the radiation process if it is lower (this is a physical lower bound on the true σ).
- Set σ to the typical value observed for NMIS measurements if it is lower.
- Increase σ in proportion to the uncertainty in its value, given the number of measurements in the template used to estimate it.

This of course resulted in $\sigma_{ij}(\tau)$ being overestimated on the average, and did not remove all of the problem values.

Past efforts had found that this technique is much improved by using a *nearest neighbor* approach. The unknown is compared to every individual measurement, *r*, in the template set to find its nearest neighbor in any template set:

$$d_{ij}^{2}(t)=\frac{\left(C_{ij}(t)-r_{ij}(t)\right)^{2}}{s_{ij}^{2}(t)}$$

This approach acknowledges that the variations in the template set are not entirely random. For example, variations in component position result in reproducible variations in the geometric distribution of the measured radiation among the detector. With a nearest neighbor approach, an unknown measured under the same conditions as one of the template measurements will find a good match. The template still needs to contain a reasonably representative set of such variations, and the template variance $\sigma_{ij}(\tau)$ is still used. This technique is still available in the software and generally identifies the unknown correctly, but this method has not been entirely satisfactory.

Finally a nearest neighbor, non-metric approach with feature weighting has been created. The unknown U is compared to each pair [m,n] of templates available. If this comparison concludes that U is more like template m than template n, a vote against template n is recorded. Template n has, in essence, been eliminated as a candidate since a more likely template has been found. A simple decision tree employing this process would be single-elimination tournament, in which U is identified with the template that beats all challengers. This simple approach occasionally fails because it considers each template decision to be definitive (single elimination). In a close decision regarding m vs n, m could be eliminated incorrectly in favor of n. If subsequent comparisons of n to other templates show decisively that n is wrong, the single elimination approach does not back up and reconsider m. Reconsideration would be futile if the same criteria were used to make all decisions: since m has already been eliminated by those criteria, the same templates that eliminate n would likely eliminate m also. But in this application different (optimal) criteria are used for every individual pair-wise template decision, so m can be accepted when n is rejected. The method currently being used successfully, but still being studied, evaluates all of the pair-wise decisions and keeps track of how many times each template was voted down. The winning template is the template with the fewest votes against it. The software also assigns a weight to each decision that reflects vote's definitiveness, described later. This technique is a form of Classification And Regression Tree (CART) [2].

Some HEU components are distinguished by all of the available $C_{ij}(\tau)$, but most are reliably distinguished only by a subset F_{mn}. In this application the Bhattacharyya Bound [2] identifies that subset for every pairing of templates. Consider two templates m and n with their individual $C_{ij}(\tau)$ values having Gaussian probability distributions [μ_m, σ_m] and [μ_n, σ_n], while the unknown component's measured value of $C_{ij}(\tau)$ is x. Assume the unknown is actually a member of template class m. If the m and n template probability distributions for $C_{ij}(\tau)$ overlap, the measurement x has some probability being closer to template n than its proper template m. The Bhattacharyya Bound is the upper bound on the error rate for this situation, when the decision criterion is simply that x will be assigned to the closest template. This bound, E, is:

$$E=\exp\left[-\frac{1}{4}\frac{\left(m_{m}-m_{n}\right)^{2}}{\left(s_{m}^{2}+s_{n}^{2}\right)}-\frac{1}{2}\ln\left[\frac{1}{2}\left(\frac{s_{m}}{s_{n}}+\frac{s_{n}}{s_{m}}\right)\right]\right]$$

Those $C_{ij}(\tau)$ which have a bound less than 1% are automatically included in F_{mn}. This is a form of feature weighting. The calculation of this bound uses the troublesome σ, but only to make a conservative judgment on whether to include a particular $C_{ij}(\tau)$ in the set of useful features. Software options to correct low σ values, described above, are still used, primarily to increase σ when a template contains too few measurements. This selection criteria works in practice.

The winning template in a pair comparison is the template which has the most unknown $C_{ij}(\tau)$ values closer to its template values (most votes). Ties are handled by lowering the Bhattacharyya Bound slightly to gather more votes. Since the decision is not based directly on the probability distribution of the measurements, as the Mahalanobis distance is, it does not depend as heavily on their being estimated very accurately; thus this method is described as "non-metric". The expected vote is nearly 100% for the proper class, given the 1% value of the bound used to pick the voting $C_{ij}(\tau)$. A 50% result is the worst possible result and contributes no information towards a final match decision. Earlier it was noted that a weight is assigned to these decisions. With N votes total and V votes for the winning template, this weight w_{nm} for the comparison of templates m and n is:

$$\frac{1}{w_{mn}} = \frac{N - V}{\sqrt{P(1-P)N}}$$

The probability distribution of votes is the binomial (two-value) distribution. The denominator term is the variance of the vote expected if the actual probability of any vote going to the correct class were P, and the numerator term is the difference between the actual vote and a 100% vote. This weight gives lower weight to results that vary from 100% and gives higher weight to results with a higher number of votes N.

These weights are calculated for each pair-wise decision made. Each template is assigned a "score", which is the lowest weight it received in this process. A template that received all votes in every comparison ($N = V$ for all decisions) receives infinite weight; a template that was soundly voted down in a high-N vote receives very little weight, regardless of how well it does in comparison to other templates. At present, this weight is used to indicate, in a general way, the confidence that the software has in the match it has found.

The nearest neighbor aspect of the method is similar to the Mahalanobis method described, but more complicated to calculate because it involves two templates and an unknown rather than one template and an unknown. The decision is made by looking at all pairings of the individual members of the two template sets. Assume two templates containing measurements n_k and m_l in templates n and m. The vote taken using n_k and m_l to represent their templates is v_{kl}. The vote associated with measurement n_k, v_k, is the worst vote recorded when comparing n_k to all of the m_l. Similarly, v_l is the worst vote recorded when comparing m_l to all n_k.

$$v_k = \min[v_{kl}] \text{ for all members } l \text{ of template } m,$$
$$v_l = \min[v_{kl}] \text{ for all members } k \text{ of template } n.$$

The winning template is the template with the best member measurement, so the final vote is:

$$v = \max[v_k, v_l] \text{ for all members } k \text{ and } l \text{ of templates } m \text{ and } n.$$

In the case that v is a tie vote, the template with the highest percentage of winning members is selected.

RESULTS

The matching method is tested against all templates and confirmation measurements as they become available. Recently the measurement time was shortened and two of the four detectors were dropped for some measurements. This was a major reason for developing a new matching technique. It also forced a switch to a moving average version of $C_{ij}(\tau)$; matching uses $C_{ij}(\tau)$ averaged over seven lags (τ). This reduces the variance in C_{ij} due to the random nature of radiation measurements.

The usual result is that the correct template wins all pair-wise template votes. In some cases, the correct template has one vote against it while the others have more votes against. To date, simple steps have been sufficient to match the correct template. Rarely, a tie is declared in which the correct template is one of two possibilities found.

As an example, a data set was made using small aluminum blocks between the ^{252}Cf source and the four detectors. Four blocks of different heights were used, designated A through D for the highest to the shortest. Eight measurements with variations in the positions of the source, detector, and blocks are the templates A through D; shorter measurements are the unknowns designated Au through Du. The software selected features only from C_{13}, the correlation between the ^{252}Cf source and the center detector. This is the best signature because the aluminum blocks' "shadow" was mostly on that detector. Fig. 1 shows these unknown signatures. The first peak is the gamma ray transmission peak, and the second is the neutrons of various energies arriving at the detector. The unknowns might appear to be easily distinguishable by the gamma response, however the variance of the gamma peak is large there as it is in other parts of the signature. For example, the A unknown (Au) is actually outside the range found in its template set at the gamma peak. Block D (Du) is most easily distinguishable. Template D was judged best for unknown Du by unanimous votes against all other templates. The result for Cu was also unanimous for template C. Blocks A and B are very similar and, in fact, practically indistinguishable in this measurement, using these signatures. Their unknowns Au and Bu are correctly identified, but the vote margin was only 7:3 over B in the case of Au, and only 6:4 over A in the case of Bu.

FUTURE WORK

In the near term, there will probably be some refinements while the template collection is being completed. Two possible areas are (1) refining the method by which signatures are selected for use, (2) replacing the voting scheme with another decision function that uses more information about the templates and unknowns, and (3) using the weight (decisiveness) of each pair-wise template decision more fully in the final decision. However, any such refinements that reintroduce template metrics into the analysis will be attempted only after an additional analysis of the templates and confirmation measurements available.

The matching software calculates a wide variety of values from the data that could be used with $C_{ij}(\tau)$ to identify a HEU component. Most of these have not been tried with the latest matching method, but we suspect they will help identification.

In the longer term, we will draw on research being done to retrieve the component's physical attributes from the measurements. We will apply what is learned there to matching.

REFERENCES

1. J. T. Mihalczo, J. A. Mullens, J. K. Mattingly, and T. E. Valentine, "Physical Description of Nuclear Materials Identification System (NMIS) Signatures," *Nuclear Instruments & Methods in Physics Research* A, 450 (2000) 531-555.
2. R. O. Duda, P. E. Hart, and D. G. Stork, Pattern Classification, Second Edition, John Wiley & Sons, New York, 2001.