

Cooperative Verification Using Radiography Behind an Information Barrier

Charles Little, Tom Weber, Chris Wilson

Sandia National Laboratories

SAND2015-4562 C

Abstract

Sandia National Laboratories has been conducting research into utilizing radiography, combined with automated image processing algorithms, to create a novel method of non-invasive verification. In many treaty verification scenarios, inspectors must verify the authenticity or identity of items that contain sensitive features. While radiography is a powerful inspection tool, it also reveals a great deal of detail about an item that may not be allowed by a verification agreement. Automation of the image processing task enables use of an information barrier, giving inspectors confirmation that an inspected item matches a previous measurement or agreed template while protecting sensitive information about the item.

Our technique utilizes feature matching in radiographic images of complex items. The SURF (Speeded Up Robust Features) method is used to extract features from the images. FLANN (Fast Learning Artificial Neural Network) is used in the matching process. The feature list becomes the template. The SURF features are somewhat rotation, scale, and translation invariant, which means the reference and target images need not be taken from the exact same position for the source and film, making data collection easier. A significant discovery is that we can discard the position information of the features and still perform the matching adequately. With no position information, geometry cannot be recovered; we believe it is impossible to reconstruct the image in this case, creating an irreversible transform that creates non-sensitive feature lists, or templates. This method is analogous to using a paper shredder to prevent reconstruction of an original while still being able to match features from the individual shredded pieces.

Results of these image processing techniques on radiography simulations are promising, showing high correlation between features from identical items, even at slightly different measurement angles. Items not matching the original have significantly lower correlation with the feature set, enabling an automated decision process.

We provide examples and results from complex electro-mechanical systems to demonstrate the effectiveness of this technique in the automatic verification of such items, and a path forward to the creation of a complete verification system with an information barrier.

Keywords: radiography, feature matching, information barrier

1 Introduction

This paper introduces a new approach for utilizing radiography behind an information barrier to verify the authenticity of an item presented for inspection, while keeping the item, and the radiography image of that item, shielded from an inspector's view. In many treaty verification scenarios, inspectors must verify the authenticity or identity of items that contain sensitive features. While radiography is a powerful inspection tool that can definitively identify an object as matching, it also reveals a great deal of detail about an item that may not be allowed by a verification agreement. Automation of the image processing task enables use of an information barrier, giving an inspector confirmation that an inspected item matches a previous measurement or agreed template while protecting sensitive information about the item.

Confirming or validating an object is often done with template matching data from the test object to reference data. For this scenario, we assume that the test object can be characterized by internal construction, which we wish to examine without dismantling. Imaging with radiography is non-invasive, non-contact, relatively fast, and rich in information. It is a good tool for verification of declarations (i.e. dismantlement) involving nuclear weapons or any other type of complex electromechanical system. The focus of our research is in verifying such a complex item against a previously generated template or reference set. Matching radiograph images would be a straight forward verification tool if visual inspection is permitted; radiography is commonly used for many non-destructive evaluation tasks. However, the geometry information acquired through radiography for treaty verification will likely contain highly sensitive information, which can never be revealed to an inspector. Therefore, all image processing in such a scenario must occur behind an information barrier, and thus be fully automated. Automated image verification tools for such an application have not been demonstrated.

This paper presents a technique that allows images to be matched through a template made up of features detected by open source image processing algorithms. The process to create and match templates can be automated, and therefore can be executed behind an information barrier with high confidence. Furthermore, we believe that if location information associated with image features is discarded, the technique is irreversible: the image cannot be recreated from the template. This method is analogous to using a paper shredder to prevent reconstruction of an original while still being able to match features from the individual shredded pieces. This results in templates which are non-sensitive. No sensitive information needs to be stored, reducing the complexity of the information barrier system as well as the procedures for its storage and use.

It is understood that in any actual verification agreement, creation of a template from a reference object presents its own challenges, notably the authentication of the original object. This paper does not address this challenge, which could potentially be accomplished through other types of more comprehensive (longer time frame) measurements or other means to establish provenance.

2 Technique

If a reference radiographic image was used directly as a template, simple image differencing could be used to match a test image of an unknown object against a known one for a complete match. However, this is really only feasible if the images are highly aligned. It fails quickly with even minor variations between the two snapshots. This type of alignment for radiograph requires controlled, fixtured settings, which are a particular problem for field gathered radiographs. It is also problematic when long time lags occur between the original imaging and later test imaging.

To resolve the problem of alignment we have researched feature detection and matching techniques,

commonly used in image processing, for application to radiography images. Feature detection and matching enable robust comparisons even with fairly unaligned views.

2.1 Feature detection

In computer vision, feature detection is a method to reduce the image data to a much smaller set of relevant information. This set of features is expressed as a feature vector. There are a number of popular methods to detect features, including SIFT (Scale Invariant Feature Transform), SURF (Speeded Up Robust Features), FAST (Features from Accelerated Segment Test), BRIEF (Binary Robust Independent Elementary Features), ORB (Oriented FAST and Rotated BRIEF), and others [1].

A literature search was completed to see which methods have found good use with radiographic images, and SIFT and SURF have proven useful [2]. Of note, this type of computer vision, though well studied, has been used primarily in visual images; while some work has been published using radiographic images, this seems to be exclusively for medical imaging. We found no direct references to feature matching in industrial radiography. In our application, SURF and SIFT outperformed all other methods. A distinct advantage to SIFT or SURF is by design their features are invariant to translation and rotation. This means the reference and target images need not be taken from the exact same position for the source and film, making data collection easier.

Both feature detection methods output, as a result, a list of feature artifacts that come in two parts. The first part, called *keypoints*, contains basic geometry information of the feature: the location, scale and orientation. The second part, called *keypoint descriptors*, is a vector of numbers derived from the algorithm, using local gradient data combined with a set of histograms processed at various orientations. Both SIFT and SURF produce 128 element feature vectors (See [3] for an explanation of SIFT). It is important to note that a feature vector is not simply a small section of the image, or an image patch.

2.2 Feature Matching

Our process compares a current radiographic image, called the *test* image, against a previous image, called a *reference* image. The reference image(s) are taken in advance, presumably in a controlled situation. The test image is taken at the time where confirmation or verification is needed, presumably in the field at a later time. There may be limited access to the object for the test image (it may be stored in a container), or difficulty in taking an image at a precise orientation or alignment. In this case, the reference set needs to be inclusive of the possible space of orientations to be expected. The goal then is to find the closest reference image, and look at the quality of the match to determine if the test object matches the object that produced the reference images. Figure 1 shows an example.




Figure 1. Test and reference imaging directions around sample target.

Two feature matching methods were implemented. The first method, which we call “one-to-one”, compares features from the test image to features from each reference image independently. Each feature from the test image is compared against every feature of the reference image. This comparison is referred to as the feature distance. Features are said to match if the distance between the feature vectors is small. The smaller the feature distance the stronger the match. The reference image that produces the minimum aggregate feature distance is the best match to the test image.

A second method for matching features, which we call “one-to-many”, compares each feature vector from the test image to all the features from all of the reference images. The minimum feature distance is determined for each test image feature. The minimum feature distances are binned by the reference image from which they were derived, producing a histogram of the minimum feature distances versus reference image. The reference image with the largest number of minimum feature distances is the best match to the test image. The aggregate minimum feature distance for all the test image features provides a metric of the strength of the match.

Our code for the feature detection and feature matching uses the OpenCV toolkit [4], which contains implementations of these algorithms in code libraries, to build the software prototype. FLANN (Fast Learning Artificial Neural Network) was used as the matching algorithm [4]. It calculates the feature distance as a vector Euclidean distance of the feature keypoint descriptors. Either method described above will produce a best match reference image to the test image. These algorithms always produce a best match whether the match is good or not; i.e. a local minimum. For this reason, the aggregate feature distance must be below a threshold to be considered a valid match.

2.3 Case Study

This case study examines the sensitivity of the image matching algorithm to ten imaging parameters: six position parameters (x, y, z rotation and x, y, z translation), two exposure parameters (energy and dwell time), and two image plane parameters (noise and focus). Each parameter was tested independently. The study was conducted using the solid model of a notional electronic timer assembly and simulated radiography images. Figure 2 shows the model as well as two radiographic images of the sample target object. The object has approximate dimensions of 10 x 4 x 2 cm.

The study used the Sandia SimXray tool from the XrayToolKit (XTK) [6] package to generate simulated radiographs for testing and generation of reference sets. The simulator can mimic several X-Ray sources, as well as films likely to be used in the field. The X-Ray source is a Betatron2, with a nominal energy of 2000 KeV and dwell time of 1000. The film is NexRay MMX with an image size of 35 by 42 cm and 2024 by 2443 pixels. Other nominal parameters include a 2 m standoff distance of the source and a 0.2 m standoff distance of the film from the center of model.




Figure 2. Sample radiograph of electronic timer model from front and from top.

Figure 3 shows an example of a test to reference image match. The images seen in the background are the test image on the left, and one of many reference images on the right. The test image would be a recent image of a target item, in this case the electronic timer model seen above. The reference image is of one the stored reference images of that model. These are not the same image; they were taken at different orientations. The yellow circles are highlighted features that were detected in each image (in this case by the SURF method). The blue lines are the pairs of features that matched. Note the upper image pair has few matches, while the lower pair has many matches.




Figure 3. Features and matches in test to reference images

In this sensitivity study, we have one test image and many reference images, each of the same object but from a different position or other parameter change. The test image was evaluated against the reference image set using both the one-to-many and the one-to-one methods.

The one-to-many method tends to produce a sharper indication of the best match image while the one-to-one method provides a better indication of the sensitivity of the image match to the study parameter. The SURF algorithm was used for all feature extraction. The geometry elements (keypoints) to all feature vectors were discarded prior to performing any matching. Our results are presented below.

Rotation about X, Y, and Z Axis - The reference images were generated at every 1 degree about the each axis. The test image was taken at a rotation of 45.5 degrees. Figure 4a and 4b show the results for rotation about the Z axis for one-to-many and one-to-one matching. Figure 4a shows a strong preference for the reference images adjacent to the test image orientation. Figure 4b shows a strong preference for the same reference images but also indicates the reference images several degrees away from the optimum on each side would still be better matches to the test image compared to anything further away. The results for rotation about the Y axis are very similar the Z axis. Since the SURF feature extraction algorithm is somewhat rotation invariant, the results for rotation about the optical axis of the imaging system (rotation about the X axis) show a broad consistently low feature distance (good match) across all of the reference images without a clear best match image preference (see figures 4c and 4d).




Figure 4a. Rotation about Z, one-to-many match.




Figure 4b. Rotation about Z, one-to-one match.




Figure 4c. Rotation about X, one-to-many match.




Figure 4d. Rotation about X, one-to-one match.

Translation along X, Y, and Z Axis - The reference images were generated at every 0.005 meters from an offset along the indicated axis from -0.15 to 0.15 meters. The test image was taken at an offset of 0.057 meters along the test axis. Figures 5a and 5b show the results for translation along the Z axis for one-to-many and one-to-one matching. Again the one-to-many method shows a strong preference for the reference images adjacent to the test image orientation. Figure 5b shows a strong preference for the same reference images but also indicates the reference images several centimeters away from the optimum on each side would still be better matches to the test image over anything further away. The results for translation about the Y axis are very similar the Z axis. Translation along the optical axis of the imaging system produces a zoom or scaling effect in the image. The SURF algorithm is scale invariant. The results for translation along the X axis are similar those for rotation about the same axis, with one-to-one results showing a broad consistently low feature distance (good match) across all of the reference images. There is a somewhat stronger preference for the nearest reference images than in the rotation case.




Figure 5a. Translation in Z, one-to-many match.




Figure 5b. Translation in Z, one-to-one match.

Energy and Dwell Time (Exposure) – For the energy test reference images were generated in 50 keV increments, from 100 to 4,000 keV. The test image was generated at 2,062 keV. We saw very little impact on image matching (aggregate feature distances) from 1400 to 3000keV (see figures 6a and 6b). The source energy level should be controllable to a much tighter tolerance than this in a real system.

The reference images for the dwell time test were generated in 50 second increments, from 200 to 4,000 seconds. The test image was taken at 1,062 seconds. One-to-many and one-to-one results for Dwell Time tests are shown in Figures 6c and 6d respectively. The one-to-one results in Figure 6d indicate a broad minimum in aggregate feature distance between 800 and 1300 seconds of dwell time. In practice the Dwell Time should be controllable to a much tighter tolerance than this.




Figure 6a. Energy, one-to-many match.




Figure 6b. Energy, one-to-one match.




Figure 6c. Dwell Time, one-to-many match.




Figure 6d. Dwell Time, one-to-one match.

Noise and Focus – The Noise and Focus tests were handled differently than the previous parameters. For noise, reference images were generated by adding an increasing uniformly distributed random noise value to the test image. In this case, the noise range is increased by one integer value per step. 200 steps were taken for the reference set. The one-to-many and one-to-one match results are shown in Figure 7a and 7b respectively. The matching algorithm appears to be tolerant of noise out to about 22 or 23 steps or a noise level of about $4 \frac{1}{2}$ bits.




Figure 7a. Noise, one-to-many match.




Figure 7b. Noise, one-to-one match.

Focus is similar to the noise problem. The reference images are generated from the probe image by adding an increasing amount of Gaussian blur. At each step the blur kernel size is increased to the next

odd square kernel size, i.e. 1x1, 3x3, 5x5, etc. 100 steps were taken for the reference set. Results of the Focus study are presented in Figures 8a and 8b. It appears that for this case the matching algorithm can tolerate about 5 or 6 steps of blur (9x9 or 11x11 kernel size). This is a significant amount of blur especially considering that the test object only spans a few hundred pixels in each dimension. In our test we were comparing an unblurred test image to blurred reference images so we are really measuring the effect of the difference in blur between test and reference image. In practice the blur in a radiograph is determined by the source to object to film distances and should be controllable to limit the difference in blur from test to reference image.




Figure 8a. Focus, one-to-many match.




Figure 8b. Focus, one-to-one match.

3 Information Barrier

Sandia National Labs has significant experience working with information barriers in verification applications. The TRIS and TRADS systems involved the protection of sensitive information both for storage and processing behind an information barrier [7][8][9]. Design of these systems required addressing red/black separation issues as well as hardware and software authentication. These systems processed fairly limited one-dimensional spectral data using compact matching algorithms and simple hardware. It is assumed that both the radiographic images of a reference object and the images taken for verification contain sensitive information that an inspector may not access. The storage and processing of 2D radiographic imagery data complicates both the hardware and software requirements within the information barrier significantly.

In the verification techniques described in this paper, both the reference images and the test image used for verification are reduced to feature vector sets. This reduces the volume of data required to represent the images and reduces the possibility of casual unauthorized observation of the data. The algorithms used to both extract the feature vectors and match feature sets are open source and will be transparent to all treaty partners. This again reduces the complexity of the information barrier structure and validation process.

The SIFT/SURF feature extraction algorithms produce separable geometry and non-geometry feature vector elements: the keypoints (geometry), and the feature keypoint descriptors (non-geometry). We have demonstrated that we can accurately match images without using the geometry elements of the feature vector. We hypothesize and are in the process of proving that if the geometry elements of the SIFT/SURF feature vectors are discarded, these feature extraction algorithms constitute an irreversible transform from the image space to the feature vector space. Thus the original image and its sensitive geometry cannot be reconstructed from the feature vector set. A literature search revealed there has been some work [10][11] on reconstructing images from SIFT/SURF feature vector sets. Both methods

referenced relied on the geometry data to accomplish their results. It is debatable that the results achieved would reveal accurate geometry or sensitive information. We have found no information suggesting that reconstruction would be possible without the geometry data. If this is true, it would make the feature vector data non-sensitive and would allow the storage of both reference template data and test image feature data, as well as the feature matching processing, to take place outside the information barrier. The information barrier system could be reduced to a system that captured a sensitive radiographic image and extracted and exported non-sensitive feature vectors. No image or other data storage, other than volatile memory used for processing, would reside inside the information barrier. Issues of how to move the information from inside to outside the information barrier have been addressed in previously developed systems such as TRIS and TRADS. The automated processing and reduced storage requirements of the feature vector matching approach is well suited for an information barrier implementation.

4 Conclusion

We have successfully used radiography images with a feature matching method to identify a test object against a reference object. The image comparison was accomplished using SURF feature detection data as the template. Once the radiograph is taken, the task of creating the test feature set and matching it against a reference template set can be completely automated. Therefore, there is no need to save the radiographic images or present them to an inspector; thus the applicability to an information barrier, because the processing can happen without user input or visibility into the actual data. We believe this new method of radiography image comparison could be of use for verification of sensitive items which must be shielded from visual inspection.

A significant discovery is that we can discard the image position information part of the feature detection and still perform the matching adequately to identify the closest reference image to the target image. With no position information, geometry cannot be recovered; in this case we believe it is impossible to reconstruct the image, creating an irreversible transform that produces non-sensitive templates consisting of feature lists. This reduces the processing of sensitive data to the extraction of features from a temporary radiographic image and eliminates the need to store sensitive images or templates, enabling complete transparency of the feature matching process.

The data shown above was generated using a simulated radiography system, with the input being a CAD model of a notional electronic timer. Results of these image processing techniques on radiography simulations are promising, showing high correlation between features from identical objects, even at slightly different measurement angles. Objects not matching the original have significantly lower correlation with the feature set, enabling an automated decision process.

We have also matched real radiography images against our CAD model based reference sets and verified good performance. This is important, since it may be difficult to obtain full reference sets by real-world radiography due to cost, safety, or security constraints. Instead, reference feature sets can be constructed through simulation only using CAD models of the real world objects. Building reference feature sets from simulated radiography based on CAD models could provide an alternative to taking large numbers of radiographic images to generate reference feature sets.

We have also analyzed the sensitivity of feature matching for radiography images that are not an exact match. We looked at the positional variables; rotation about the Y and Z axes (elevation and azimuth), and translation about X, Y, and Z axes. We found that translation is fairly insensitive, meaning we can position the object within about 10 cm without much degradation. Rotation is much more sensitive, but

as can be seen from Figure 4 above, one or two degrees of rotation has little effect on feature correlation, and good correlation can be found even with up to 10 degrees of rotation difference between test and reference images. We also tested other radiography variables: energy, dwell-time (exposure), focus, and image noise. These are all fairly insensitive to change, and can reasonably be expected to be controlled in the setup of the radiography system.

There are several other applications or extensions that may be areas for future work. Application as a unique identifier for chain-of-custody purposes could potentially identify unique manufacturing differences rather than the global similarities of type. Another application could be to verify a range of different object types by substituting template sets. If there is a need to identify the type of object, all the necessary templates could be installed on the same system. Finally, it may be possible to automate feature recognition for diagnostics or emergency response. Identifying missing/removed or damaged/displaced components in an object might be accomplished by including geometry information and identifying non-matching regions compared to a reference.

References

- [1] http://docs.opencv.org/modules/features2d/doc/feature_detection_and_description.html
- [2] L. Schmidt-Hackenberg, M. R. Yosuefi, T. M. Breuel, "Visual Cortex Inspired Features for Object Detection in X-Ray Images," International Conference on Pattern Recognition November 2012.
- [3] David G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints", International Journal of Computer Vision, 2004.
- [4] <http://docs.opencv.org/modules/core/doc/intro.html>
- [5] Marius Muja and David G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," International Conference on Computer Vision Theory and Applications (VISAPP), Lisbon, Portugal, Feb 2009.
- [6] http://www.sandia.gov/research/robotics/assets/documents/XTK_Handout_Final.pdf
- [7] Seager, Mitchell, et. al., "Trusted Radiation Identification System", Proceedings of the 42nd INMM Annual Meeting, July 2001.
- [8] Merkle, Weber, et. al., "Next Generation Trusted Radiation Identification System", Proceedings of the 51st INMM Annual Meeting, July 2010.
- [9] Tolk, Mitchell, "Trusted Radiation Attribute Demonstration System", Proceedings of the 41st INMM Annual Meeting, July 2000.
- [10] P. Weinzaepfel, H. Jegou, and P. Perez, "Reconstructing an Image From Its Local Descriptors," in Computer Vision and Pattern Recognition, Colorado Springs, USA, Jun. 2011.
- [11] https://stacks.stanford.edu/file/druid:my512gb2187/Daneshi_Guo_Image_Reconstruction_from_Descriptors.pdf

An Information Barrier for Radiography Image Verification


SL16-RadInfoBarrier-PD2Nc

Sandia National Laboratories: Tom Weber (PI), Maikael Thomas, Charles Little, Chris Wilson

Pantex: Allie Williams and Monte Epperson.




DATA ACQUISITION




The demonstration system...

...consists of a COTS BeagleBone Green to create feature sets and perform matching. A custom printed circuit board connects to the BeagleBone and provides system interface cues.



The Reference Feature Set is digitally signed to provide information assurance



COMPARE

Test image feature set against reference image feature set

MATCH

If the test data set matches any one of the reference feature sets, the test item is authentic.




Image 01

45° Camera Angle








Image 02

60° Camera Angle



Image 03

75° Camera Angle



NO MATCH

If none of the reference feature sets correspond to the test data, the test item is not authentic.




REFERENCE

Multiple images, individual angles on item. Item is shielded from view


TEST

Single image from single angle. Item is shielded from view



Item 01

Item 02



Test data feature set is generated



8th INMM/ESARDA Joint Workshop: Position Paper

Arms Control: Balancing transparency vs. protection of sensitive information

Arms Control: Novel Technologies for Arms Control Verification

Cooperative Verification Using Radiography Behind an Information Barrier

Charles Little, Tom Weber, Chris Wilson: *Sandia National Laboratories, Albuquerque, NM*

SAND2015-6331C

Position Statement

Treaty inspectors must verify the authenticity or identity of items that contain sensitive features without actually revealing the sensitive data. Sandia National Laboratories has developed a technology utilizing radiography, combined with automated image processing algorithms, to create a novel method of verification that does not process or store sensitive information.

Radiography is a powerful inspection tool. Direct image comparisons can be used to verify the existence or absence of parts. It can also be used to authenticate that a current part matches a previously imaged part. However, it also reveals a great deal of detail about an item, some of which is sensitive data that may not be revealed by a verification agreement.

Our technique utilizes feature matching in radiographic images of complex items. The SURF (Speeded Up Robust Features) method is used to extract features from the images. FLANN (Fast Learning Artificial Neural Network) is used in the matching process. The feature list becomes the template. The SURF features are somewhat rotation, scale, and translation invariant, which means the reference and target images need not be taken from the exact same position for the source and film, making data collection easier. A significant discovery is that we can discard the position, size, and orientation information of the features and still perform the matching adequately. Without this information, geometry cannot be recovered; we believe it is impossible to reconstruct the image in this case, creating an irreversible transform that creates non-sensitive feature lists, or templates. This method is analogous to using a paper shredder to prevent reconstruction of an original while still being able to match features from the individual shredded pieces.

We propose a non-recording radiography system, combined with our image processing technologies to create non-sensitive template data, which can be used behind an information barrier. The matching process, which uses non-sensitive data, can be completed outside the information barrier; this enables a trusted verification system, reducing the complexity of the system within the information barrier.

Results of these image processing techniques on radiography simulations are promising, showing high correlation between features from identical items, even at slightly different measurement angles. Items not matching the original have significantly lower correlation with the feature set, enabling an automated decision process. These results are shown in figures 1 and 2 below.




Figure 1: Radiograph to non-sensitive feature template




Figure 2: Feature matching for verification

We have also matched real radiography images against our CAD model based reference sets and verified good performance. This is important, since it may be difficult to obtain full reference sets by real-world radiography due to cost, safety, or security constraints. Instead, reference feature sets can be constructed through simulation only using CAD models of the real world objects. Building reference feature sets from simulated radiography based on CAD models could provide an alternative to taking large numbers of radiographic images to generate reference feature sets.