

This is a peer-reviewed, accepted author manuscript of the following article: Devereux, M. G., Murray, P., & West, G. (Accepted/In press). Automated object detection for visual inspection of nuclear reactor cores. Nuclear Technology.

# Automated Object Detection for Visual Inspection of Nuclear Reactor Cores

Michael G Devereux,<sup>\*,a</sup> Paul Murray,<sup>a</sup> and Graeme West<sup>a</sup>

*<sup>a</sup>Department of Electronic and Electrical Engineering, University of Strathclyde,  
204 George Street, Glasgow, G1 1XW, U.K.*

\*Email: [michael.devereux@strath.ac.uk](mailto:michael.devereux@strath.ac.uk)

Number of pages: 28  
Number of tables: 2  
Number of figures: 14

## Abstract

Remote Visual Inspection (RVI) is a common approach to understanding the health of key components and sub-structures within nuclear power plants, particularly in difficult to access and high dosage areas. Interpretation of inspection footage is a manual intensive procedure and challenges arise in localizing and dimensioning defects directly from a video feed, which may be subject to uncertainty from a range of sources such as lens distortion, non-uniform lighting and lack of depth from a monocular camera system. A common approach to addressing these issues is to develop a scaling factor based on identifying a reference object of known dimensions in the image and using this to size regions of interest. Manual, accurate identification of these reference objects is onerous, time consuming and prone to variation across different human experts, therefore, robust identification of suitable reference objects in an automated, reliable and repeatable manner is of significant value. In this paper we evaluate two approaches for the automated detection of reference objects in the inspection of graphite cores in the UK's fleet of Advanced Gas-Cooled Reactors (AGR). The first method is a multistep approach using tools from mathematical morphology. The approach uses a genetic algorithm to “grow” suitable structuring elements, refine the order of operations and remove operations proposed by the human designer that have a negative impact on performance. The second approach uses semantic segmentation, a technique which is normally applied to scene labeling in computer vision applications, applied to produce a binary mask, separating the reference object from the background. We show that this second method performs significantly better than the mathematical morphology approach when applied to the identification of brick interface keyways in AGR inspection images. Though improved in terms of accuracy, it is recognized that a greater initial effort is required to train the approach, and as it utilizes black-box neural network approaches, the greater transparency offered by the mathematical morphology approach is lost. While explicability of techniques is often a highly desirable characteristic of automated analysis techniques applied to health assessment within nuclear power plants, the results of the reference object detection can be made explicit to the end user, ensuring that the human analyst is retained within the decision-making process thus mitigating the need for transparency.

**Keywords** — Transfer learning, Hit-or-miss transform, Nuclear power plant

## I. INTRODUCTION

An important activity required for the continued safe operation of nuclear power plants is a detailed understanding of the condition of the reactor internals. Inspections carried out during routine planned periodic outages are a primary means of understanding the health of the reactor internals. The type and frequency of inspections, vary between different countries and reactor designs, but the underlying approach is common. This approach is to gather appropriate and representative data on key components and manually evaluate and report the condition of these components for return to service of the plant. This evaluation tends to encompass the two related activities of detection of predefined indicators of health and the subsequent quantification of the severity of these indicators. Traditionally this evaluation is a manually intensive and laborious process undertaken by a domain specialist drawing upon years of operational experience. The types of inspection data usually gathered and interpreted are either videos and images from remote visual inspection equipment or, where volumetric assessment is required, ultrasonic scans, or similar.

Recent advances in processing and analyzing condition monitoring data have shown considerable benefits which can be gained through application of data-driven analysis techniques [1, 2]. However, when decisions are being made to return the reactor to service there is burden of explicability, or justification to regulator, which poses challenges for adopting many of the black-box data driven techniques.

In this paper we focus on the application of tools from machine learning, mathematical morphology and mathematical optimization to support visual inspection, and in particular the reliable and repeatable detection of reference objects with a known shape and physical dimensions observed in inspection videos. Importantly, the application of these techniques is framed as decision support, rather than a fully autonomous approaches, which means the human analyst is still retained in the decision making process. We compare and contrast two different state-of-the-art approaches to object detection, one drawn from mathematical morphology coupled with optimization and one from semantic segmentation using deep neural networks (DNN) in terms of their accuracy and explicability. A case study targeting visual inspection of AGR cores is provided, with a focus on detecting keyways as reference objects. The remainder of the paper is structured as follows: Section [I.A](#) provides a review of state-of-the-art in automated interpretation of reactor internal inspection videos, Section [I.B](#) provides a brief background of AGR core inspection and Section [I.C](#) details the

nature of keyways in inspection images. Section II introduces the novel methods for detecting keyways in AGR footage and outlines a method for very accurately measuring the dimensions of the keyways. Section III compares and contrasts the techniques and discusses their application within a remote visual inspection context with Section IV finishing with the conclusion of the paper that for detection of keyways in AGR footage that semantic segmentation yield excellent results while the optimized morphology yields very good results and is much better in terms of explicability.

### I.A. Related Work

The components, degradation mechanisms and types of data gathered varies greatly across different reactor types. In the CANDU (CANada Deuterium Uranium) reactor design, the zirconium alloy pressure tubes which form the fuel channels are inspected ultrasonically, where a multi-headed ultrasonic probe is used to gather precise depth measurements over the full inner surface of the tube, and any individual depth measurement which breaches a pre-determined threshold needs to be assessed and evaluated under fitness for service criteria before restart of the station [3, 4]. Reactors of the VVER design are also inspected using ultrasonic equipment. The examination of the vessel inner surface combines ultrasonic data with eddy current analysis to gain information about the condition of reactor welds and perform clad surface examinations [5]. Boiling Water Reactor (BWR) and Pressurized Water Reactor (PWR) are thick-walled pressure vessels where routine inspection of the welds is required [6]. Inspections are usually performed underwater in the pressure vessel and inspections are performed using remotely operated submersibles [7, 8, 9] with the footage manually assessed. In Advanced Gas-cooled Reactor (AGR) stations, the graphite fuel channels are inspected through a combination of visual inspection and accurate physical channel bore, tilt and ovality measurements [10, 1, 11].

This paper is focussed on extracting useful information from remote visual inspection data and there has already been some work undertaken in the application of automated analysis of image data in the nuclear industry. To date these have focussed on supporting the detection of defects and regions of interest such as cracks. Ref. [12] focus on the automated detection of cracks on steel plates using a three stage approach consisting of segmentation, filtering out areas of weld and anomaly detection. This approach performs much better on a limited dataset when compared with the morphology approach for crack detection proposed by [13]. Ref. [14] use advances in deep

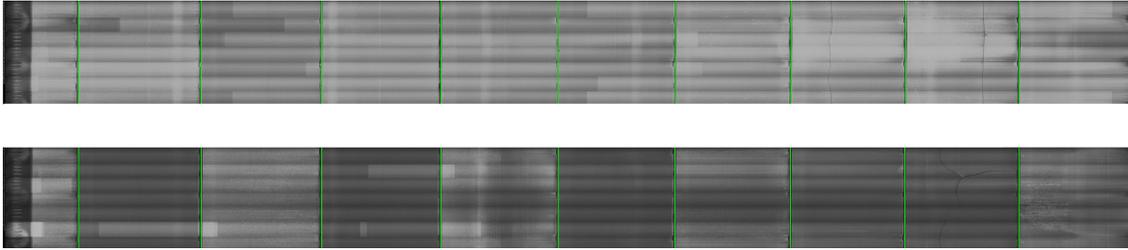


Fig. 1. Example Chanoramas. The brick interfaces are denoted by green vertical lines.

neural networks by retraining existing convolutional neural networks to classify image patches of both real and laser machined cracks on steel. This approach was compared with a popular automated crack detection toolbox CrackIT [15] and performed significantly better. Ref. [16] refine these techniques further to classify individual pixels in image patches as crack/not crack by building on their deep neural network approach and by aggregating the same region of multiple frames to classify each pixel. Ref. [17] developed a framework a deep learning framework, based on a convolutional neural network (CNN) and a Naive Bayes data fusion scheme to detect cracks in a synthetic dataset that represents internals of nuclear power plant components. This approach is compared with a prior framework developed by the research team also designed to detect cracks using local binary patterns, support vector machines and data fusion [18].

Ref. [19] developed a probe for inspecting 3-6 inch diameter pipes for application in inspection of nuclear reprocessing pipework. This work focussed on a novel means of visualizing the data (generating the images) rather than supporting the assessment of the generated images. Ref. [11] present a system ASIST (Automated Software Image Stitching Tool) for automatically converting in core inspection videos into chanoramas (fuel channel panoramas). Again this approach assists in the visualization of the data rather than its interpretation. Ref. [20] propose a method for the automated detection of trepanned holes in the resultant images produced by ASIST with the aim of calculating scaling factors to allow the accurate sizing of defects such as cracks. The paper also describes another system for the automated detection of cracks and a novel means of visualizing crack likelihood in fuel channel images.

Each of these automated approaches aim to aid a skilled operator in the assessment of inspection images by automatically highlighting areas of interest. These works do not consider how to use these areas of interest to quantify the severity of any defects observed. The dimensions of a defect observed in inspection images are unknown so it is important to measure the dimensions of

these objects. This can be achieved in two ways; accurate camera calibration or using scaling factors calculated from known physical objects in the images. Camera calibration is used to estimate the parameters of an optical system to correct any distortions in the recorded images which allows the measurement of the size of an object in an image in real world units. Camera calibration is achieved by recording a known standard such as a checker board at multiple different poses to allow camera parameters to be estimated [21]. Camera calibration is not possible for every optical system and in cases such as these it is also possible to make accurate real world measurements of objects of unknown dimension such as cracks using reference objects in the images where the physical dimensions are known. This is achieved by calculating a scaling factor using this reference object which equates pixel measurements with real world metric measurements. The system proposed by [20] calculates scaling factors using trepanned holes but trepanned holes are only present in a very small minority of AGR fuel channels. In this work we consider the automated detection of keyways which appear in every fuel channel of 5 of the 7 AGR stations in the UK. Keyways also have a known dimension which make them very suitable for calculating scaling factors.

### **I.B. Remote Visual Inspection**

Visual inspections are routinely performed by inserting a radiation hardened camera mounted on a specialized tool [10] and recording the video feed from the camera. This camera is mounted so that the field of view is looking down the fuel channels but it is possible to record the inside wall of the fuel channel using a mirror that is part of the inspection tool. The position of this mirror is controlled by a motor so it is possible to view either down the fuel channel or the side walls as required. Routine inspection videos consist of 6 overlapping scans with a field of view of 70 degrees. Scans are recorded at 6 different orientations and for each scan the camera tool is rotated 60 degrees which allows the entire inner surface to be recorded. A chanorama is a full montage spanning the inner surface of a fuel channel created from these scans using a purpose built piece of software known as ASIST (two examples shown in Figure 1). ASIST works by automatically decomposing the frames from inspection videos, stitching each scan into an individual strip and then stitching every scan strip together to form a chanorama.

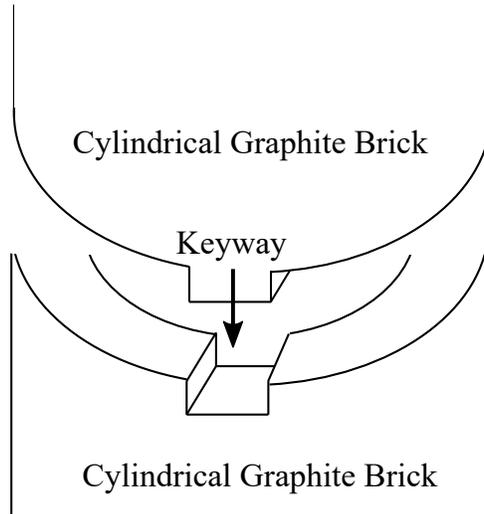


Fig. 2. Keyway schematic.

### I.C. Keyways

The graphite bricks used to construct the fuel channels are joined together at the brick interfaces using four end-face keys on the base of each brick which interlock with keyways on the top of the brick underneath (see Figure 2). These distinct channel features can be seen in the chanoramas, see Figure 3. The shape and size of the interlocking keys and keyways at each station are known and their appearance is reasonably consistent within a single chanorama. However, keyway appearance can vary significantly from chanorama to chanorama and especially from station to station due to aging effects and effects from the inspection tool recording the footage. Furthermore, the chanoramas generated using inspection footage are not perfect due to distortions and variations in lighting introduced by the camera and imaging system. These distortions present as the keyways appearing wider when observed in the left or right hand regions of the chanorama strips compared with keyways observed in the centre of chanorama strips. These distortions can be negated by recording all features of interest in the centre of the frame which is the usual procedure when cracks are observed. The width of the keyways is accurately known for all reactors at the start of life allowing the keyways to serve as accurate reference objects for equating real world distances to pixel measurements.

In this paper we focus on the problem of automatically detecting keyway locations and then automatically measuring their pixel dimensions regardless of which station they belong to. The imperfections described above make designing a generalized keyway detection framework challenging.

There are a number of problems associated with panoramas where knowing the location and widths of every single keyway would be useful. These include

- Calculating scaling factors to measure the width of defects observed in the fuel channels during inspections as demonstrated by Ref. [20].
- Estimating the rotational drift of the camera as it transverses the fuel channel during each scan.
- Estimating errors in strip placement by considering the relative positions of keyways and their relative differences.
- Registering images from the same fuel channel recorded during different inspections.

To manually label and measure the dimensions of every keyway for each panorama would be a very time consuming process for a skilled human operator and prone to variation from different operators. Therefore, a reliable and repeatable automated means of detecting and measuring the keyways is of great benefit and can offer significant time savings.

We present two new automated methods for detecting and locating keyways. The first considers a tool from mathematical morphology known as the Hit-or-Miss Transform [22] combined with suitable pre and post-processing. Mathematical morphology has been applied to a wide range of problems and achieves considerable success in the automated detection of various different objects in images ranging from ship detection in satellite images [23] to galaxies in astronomical images [24]. It is also very suitable for this type of problem due to the fact that it natively works on pixel by pixel basis as opposed to an object basis. This framework is optimized using a genetic algorithm to provide a robust and transparent framework to automatically detect keyways. While prior research [25, 26] uses a genetic algorithm to design a single structuring element to perform an image processing operation, our system designs multiple structuring elements, orders their use and combines them in the most effective way to construct a classifier suitable for labeling individual pixels as keyway/not keyway.

The second method considers a state-of-the-art technique from deep learning known as semantic segmentation which classifies every pixel in the chanorama images. Semantic segmentation using deep neural networks has demonstrated exceptional performance in labeling natural scenes [27]. The problem considered in this work is different as we have one pixel class of interest (keyways) and we are not concerned with labeling the rest of the scene. We adapt Segnet, a semantic segmentation framework by using transfer learning to automatically label keyway pixels. Transfer learning of deep neural networks is a technique which uses a pretrained neural network and replaces the appropriate layers to allow the network to be used for different problems [28]. Semantic segmentation and Segnet in particular is still very suitable for the task of keyway detection due to the fact it also works on a pixel by pixel basis and is capable of achieving the high performance typical of deep neural network approaches. Deep learning frameworks for classification and labeling tend to provide better performance than more traditional techniques such as mathematical morphology in image processing but lack the explicability of operation that is required in the tightly regulated nuclear industry. Furthermore they are very computationally expensive during the training stages in comparison to morphological approaches.

We compare the performance of each method and outline the advantages and disadvantages of each method. We also propose an algorithm to very accurately measure keyway dimensions in pixels using the detection locations from the previous methods as seed points. These tools provide very fast and repeatable ways of detecting keyways and performing automated analysis compared with a human operator which can result in reduced time for inspections. While this paper focuses on a refinable object detector in the context of AGR core inspection imagery, the proposed frameworks could be adapted to detecting physical features of known shape in other types of imagery. The key contributions of this paper are two methods suitable for the accurate localization of keyways in AGR inspection images. The choice of which to use depends on whether performance or explicability are more important in practice.

## II. AUTOMATED KEYWAY DETECTION

In this section we describe the algorithms used to automatically detect keyways. To detect keyways in a general way suitable for all reactor inspection footage, a robust and repeatable technique is required which efficiently returns keyway locations. A simple method using a 1D intensity

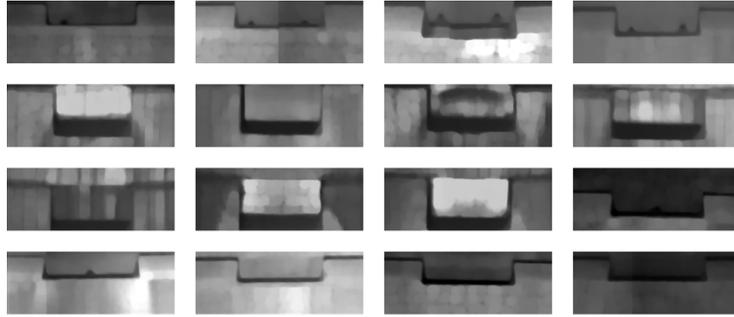


Fig. 3. Examples of keyways showing significant variation due to the images being recorded in different fuel channels or stations.

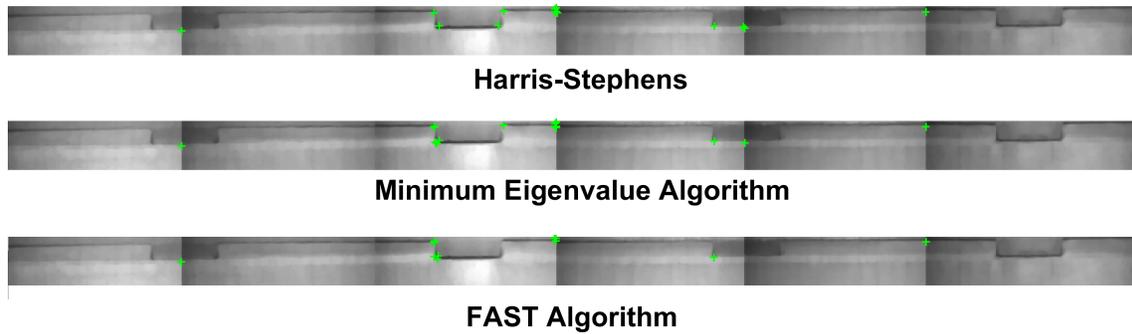


Fig. 4. Attempting to detect keyways using corner detection. The corners detected are denoted by green + symbols.

plot representing the sum of the pixel values for each row and searching for peaks corresponding to the keyway sides was considered, but was not sufficiently robust to accommodate instances where brick interfaces were not perfectly horizontal. It was also not able to discriminate between interfaces and other features from the image which share similar properties to the keyways. However, this approach is still useful to very accurately refine the exact position of the keyway sides after the keyways have been localized using other methods by considering only the local area for each keyway detection. Furthermore, due to the rectangular nature of the keyways shown in Figure 3, it is reasonable to assume that it may be possible to detect keyways using automated corner detection techniques. However this was not found to be the case. Given in Figure 4 is an example of a brick interface containing four keyway where we attempt to accurately detect the corners using the FAST algorithm [29], minimum eigenvalue algorithm [30] and Harris-Stephens algorithm [31]. This approach does not work and results in both missed keyway corners (false negatives) and a large number of false positives in undesirable locations. A large number of these false positives

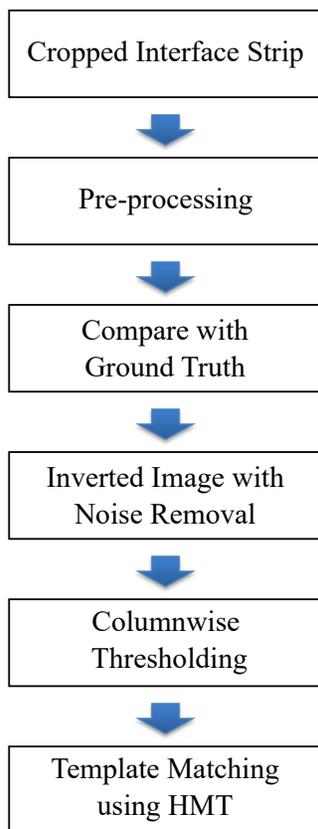


Fig. 5. Outline of keyway detection algorithm.

occur on the seams of the chanorama strips and would present a significant challenge in using corner detection methods as a means of locating keyways. Therefore, it is clear that we require a frameworks which detects the characteristic shape of the entire keyway.

### II.A. Keyway Detection Using Mathematical Morphology

We first present a multistage approach using tools from mathematical morphology to automatically detect keyways in chanorama images. It is possible to automatically detect brick interface locations in chanoramas with very high accuracy [20] and since the keyways (when present) are only seen at the interface, we need only consider these small horizontal regions of the chanorama when searching for keyways. Given in Figure 5 is a diagram outlining the overall approach.

The first step in the detection process is to calculate the complement image and remove noise from Figure 6(a) using a median filter and a morphological opening resulting in Figure 6(b). For

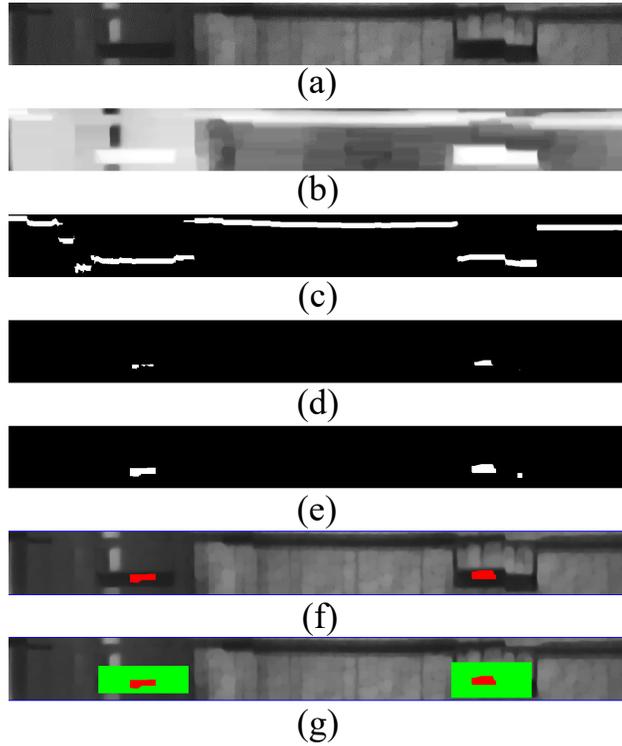


Fig. 6. Interface strip images during each stage. (a) Original image, (b) complement image with noise removed, (c) binary image with a selected number of brightest pixels in each column as foreground, (d) binary keyway detection mask after applying the HMT, (e) dilation of binary detection mask, (f) binary detection mask overlaid on original image, (g) binary detection mask compared with ground truth.

a binary image  $X$ , the morphological opening is defined as

$$O = (X \ominus b) \oplus b \quad (1)$$

where  $\ominus$  denotes erosion,  $\oplus$  denotes dilation and  $b$  is a structuring element designed for a particular application by the user. Erosion and dilation are the fundamental operations used in mathematical morphology which increase or decrease the area of foreground objects respectively in a binary or grayscale image. Erosion is defined as

$$X \ominus b = \{x | b_x \subseteq X\} \quad (2)$$

and dilation is defined as

$$X \oplus b = \{x | (\hat{b})_x \cap X \neq \emptyset\} \quad (3)$$

where  $\hat{b}$  denotes the reflection of the structuring element (SE)  $b$  [32]. In our approach, these operations are applied to the image complement before enhancing the image contrast using gamma correction. We then exploit the fact that we are seeking horizontal lines to obtain a binary image by thresholding the brightest  $k$  pixels for each column in the de-noised, contrast enhanced, image complement. To locate the keyways in the resulting binary mask, another transform from Mathematical Morphology known as the Hit-or-Miss Transform (HMT) [22, 33, 34] is proposed as a robust approach. The HMT is a template matching method which uses two complementary SEs,  $b_{fg}$  and  $b_{bg}$  where  $b_{fg} \cap b_{bg} = \emptyset$ . The HMT detects locations where the image foreground matches  $b_{fg}$  while simultaneously  $b_{bg}$  matches the image background. The HMT for the binary set  $X$ , may be defined as follows:

$$HMT_{(b_{fg}, b_{bg})}X = (X \ominus b_{fg}) \cap (X^C \ominus b_{bg}) \quad (4)$$

where  $X^C$  denotes the complement of  $X$ . Applying the HMT to the binary image obtained by thresholding, as shown in Figure 6(c), yields a binary detection mask where one valued (white) pixels represent the presence of keyways - see Figure 6(d). We can further refine this binary detection mask by applying the following post-processing rules

1. Any detected pixels which lie upon the already detected brick interface are removed;
2. The binary detection mask is dilated using a small rectangular structuring element to connect any small clusters of detection pixels which are in close proximity (left of Figure 6(d));
3. Clusters of pixels whose area is greater than  $A_{max}$  pixels based on the known physical size of keyways are removed;
4. Clusters of pixels whose area is less than  $A_{min}$  pixels are assumed to be noise and are thus removed.

These post processing steps are useful for removing false positive detections introduced by noise in the images as well as erroneous detections made at the brick interface (whose location is already

known). The final result shown in Figure 6(e) is a binary detection mask where each pixel cluster represents the location of each keyway in the image. Choosing suitable shapes and sizes for morphological structuring elements required for the HMT and postprocessing of the binary detection mask is a non-trivial task and can require considerable expertise. Furthermore, due to large variations in the appearance of keyways observed between channels and stations and the large number of parameters<sup>a</sup> within the proposed framework, it would be very difficult for a human operator to manually fix these parameters in an optimal way for general application of the algorithm to data from all stations. A genetic algorithm (GA) is an excellent choice for addressing this particular challenge due to its ability to converge to a global minimum for a wide range of problems and due to its native support for mixed integer programming [35, 36, 37]. This is particularly useful as all SE sizes correspond to integer numbers of pixels and, given a list of common shapes (disk, line, square, etc.), the genetic algorithm can also be used to determine the most suitable one to be used for a particular application. For any optimization routine an objective function is required to minimize and, in the case of the genetic algorithm, this is known as the fitness function. A suitable fitness function to maximize keyway detection performance is the negative value of the  $F_1$  score. Further details of this metric are described and presented in the results section where the same  $F_1$  score is also used to evaluate the performance of our algorithms. In practice, for the specific application of detecting keyways in AGR fuel channels it is also possible to exploit the regular placement of these features in the chanoramas to improve the performance. Doing so makes it possible to impose a regular lattice structure on the detection map in order to predict and interpolate the position of any keyways missed by the HMT as well as to remove false positive detection which do not conform to the expected structure. In this paper, however, we evaluate the more general algorithm described here without this extension to present its generality for a wider range of applications.

## II.B. Keyway Detection using Semantic Segmentation

We now consider a state-of-the-art technique originally designed for segmenting natural scenes and apply it to automatically label the pixels of a chanorama into two classes keyway or non keyway. For this work we consider Segnet developed by [27] and implemented as part of the MATLAB

---

<sup>a</sup>Examples include the dimensions of the various structuring elements and pixel areas of detection objects in binary detection masks.

Computer Vision Toolbox. Segnet was originally designed for the automated segmentation of video footage recorded from automobile camera with possible applications in self driving cars and was trained using the CamVid dataset [38] from the University of Cambridge. Segnet takes each frame and labels every pixel as one of a finite number of predetermined classes such as car, lorry, bicycle, building, tree, etc. Segnet is one type of convolutional neural network (CNN), a deep neural network specifically designed for image classification tasks. Since the publication of [39], it is now feasible to train a deep neural network from scratch to automatically classify images. Convolutional neural networks not only classify images into different classes, they also learn features (visual descriptors) that are used to classify the data. This is contrary to a more traditional image classification framework which usually operates using custom crafted features (eg. SURF [40]) to extract features for each image in a data set and then uses a separate classification stage. The main disadvantage of training a CNN from scratch is that a CNN requires enormous amounts of data, typically in the tens of thousands of instances to learn these features. Training is also very computationally expensive and can take many days on a modern GPU. However it is possible to repurpose a pretrained deep neural network using transfer learning. This is due to the fact that many of the visual descriptors calculated during training for a CNN may still be applicable in different image classification tasks. Retraining is achieved by removing specific layers from the deep neural network, replacing them with layers appropriate to the new classification task and fine tuning the network using a much smaller data set with examples of the new image classes. A more detailed introduction to deep neural networks is beyond the scope of this paper so the reader is directed to Ref. [41] for further information.

We repurpose Segnet to automatically detect keyways at the interfaces of bricks using transfer learning and the same training and testing datasets we used in the previous section to optimize the mathematical morphology approach. There are a number of challenges to overcome to allow Segnet to work with our dataset. Segnet requires image in a very specific format to work with the pre-existing neural network. It requires all images used to retrain the network to be exactly in the same format as the original network images. While this will not present much of a challenge for many applications, the images of the keyway strips are significantly larger than Segnet's training images ( $360 \times 480$  pixels) while the keyway strips are typically ( $150 \times 3000$  to  $150 \times 3500$  pixels). Two potential solutions to this problem are to either further tile the images into the data set into

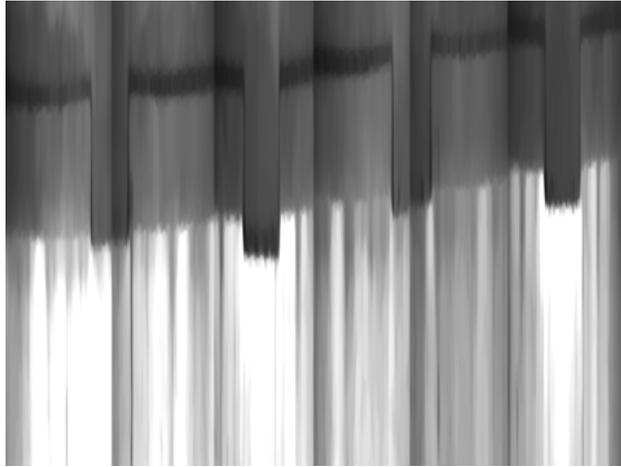


Fig. 7. Keyway strip image distorted to dimensions suitable for use with Segnet.

a suitable size or resize the images. Given that tiling could possibly break a keyway into multiple images and resizing is still required as the image height is less than half that required by Segnet, the best approach is to resize the images to match Segnet’s requirements. This results in a significant distortion of the keyway strip images as shown in Figure 7 but this does not degrade performance as it is the difference in features that is important and the images can be reshaped as required after classification is performed. This is demonstrated in Section III.B.

### II.C. Keyway Regrowing

Thus far we have presented two possible approaches to automatically locate keyways in chanorama images. In Section III we will demonstrate that while both approaches locate the keyway with a high degree of accuracy, neither return a clean binary detection mask that could be used to ascertain the keyway’s dimensions. Using the centroids of objects in the binary detection mask generated by either algorithm as a seed point, we can accurately locate the side walls of the keyway and the top and the bottom if needed. We do this by considering each individual keyway in turn resulting in images like that shown in Figure 8. This image is then converted to a 1D signal by summing each individual column’s intensity. This results in a 1D signal with the same number of points as the input image width. The sides of the keyways appear in this signal as a sharp intensity change. To aid in the automated extraction of these locations, the signal is passed through a high pass filter to remove all but sudden changes in intensity resulting in Figure

9. The two strongest positive peaks now correspond to the locations of the sides of the keyway in the original image and can be easily detected using a standard peak detection algorithm. These positions can then be plotted on the original image as shown in shown in Figure 10.



Fig. 8. Original keyway image.

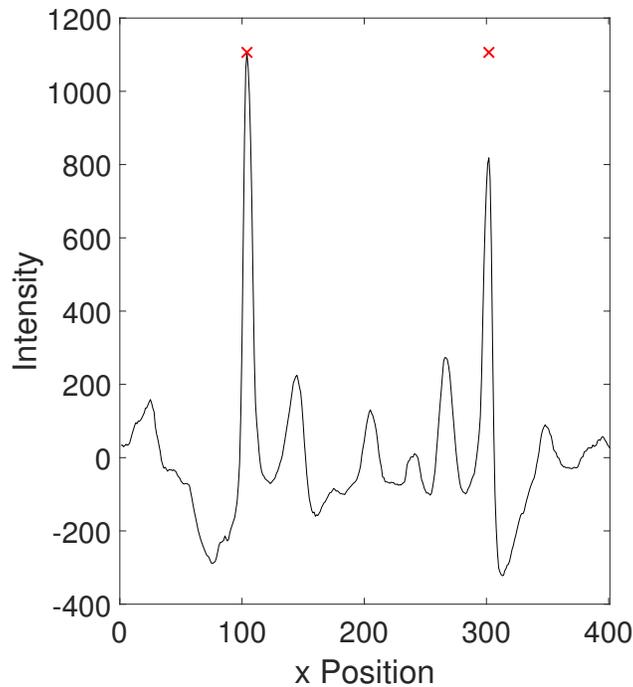


Fig. 9. Detecting the sides of a keyway by converting to a 1D signal.

### III. RESULTS

To evaluate the overall performance and test generalization of each detection framework, suitable performance metrics must be chosen. The detection of keyways in a panorama image can be considered a binary classification problem. We can label the two classes as either positive i.e. keyway present and negative i.e. keyway absent. Popular choices of performance metrics for binary classification problems include true positive rate (recall), false positive rate (fall out),



Fig. 10. The sides automatically detected are marked with a red cross.

precision, accuracy, etc. However, some of these metrics require a finite number of true negatives and, for each of the frameworks proposed herein, at the object level, there are virtually infinite true negatives. A suitable performance metric for assessing our proposed techniques is therefore the  $F_\beta$  Score. This is the weighted harmonic mean of precision and recall and is defined as

$$F_\beta = (1 + \beta^2) \frac{\text{Precision} \times \text{Recall}}{(\beta^2 \times \text{Precision}) + \text{Recall}} \quad (5)$$

where Precision and Recall are defined by

$$\text{Precision} = \frac{\Sigma \text{ True Positives}}{\Sigma \text{ True Positives} + \Sigma \text{ False Positives}} \quad (6)$$

and

$$\text{Recall} = \frac{\Sigma \text{ True Positives}}{\Sigma \text{ True Positives} + \Sigma \text{ False Negatives}} \quad (7)$$

The most commonly used variation is the  $F_1$  score in which precision and recall are weighted equally ( $\beta = 1$ ).

To evaluate statistical performance on this particular problem we consider a data set of 19 chanoramas in total. While this appears to be a small dataset, each chanorama contains 9 brick interfaces with at least 4 keyways so the overall dataset contains in excess of 600 keyways. In some cases, chanoramas have 5 keyways per brick layer due to one scan being repeated in what is known as a 7 strip chanorama<sup>b</sup>. Some of these additional keyways are only partially visible so they

---

<sup>b</sup>Chanoramas are typically constructed from 6 scans but the final result may contain 7 scan images with the first scan repeated after the last scan. The purpose of this is to ensure that all regions are completely imaged without breaks in features. This results in some parts of the channel being repeated and when keyways lie on the edge of the first image strip, they can be repeated.

TABLE I  
HMT keyway detection results.

		Actual	
		Positive	Negative
Prediction	Positive	TP = 308	FP = 25
	Negative	FN = 43	TN: NA

are considered unique for the purposes of evaluating the detection frameworks but not considered when evaluating the keyway dimension algorithm.

### III.A. Mathematical Morphology

To evaluate the performance of our proposed methods, objects in the final detection mask produced by the proposed algorithm are compared with a manually created ground truth binary mask. For a detection to be considered valid, the number of overlapping pixels in the automatically produced detection mask and the ground truth mask must exceed a predefined threshold<sup>c</sup>. Parameters (SE sizes, shapes, thresholds) were optimized using 10 of the chanoramas (4 different stations) from the data set of 19 chanoramas. We evaluated the algorithm’s performance on a test set of the remaining 9 chanoramas (5 different stations including the 4 in the training set) containing 351 keyways. While many machine learning approaches typically split the data 70/30 for training and testing, this distribution was chosen to ensure multiple chanoramas from each station were contained in both the testing and training sets. Due to the large number of keyways in each chanorama, excellent performance can still be achieved with this relatively small training data set. The results from evaluating the proposed algorithm on the 9 chanoramas is summarized in Table I. In this table, TP denotes true positives, FP denoted false positives, FN denotes false negatives and TN denotes true negatives. The number of true negatives is not applicable (NA) for this particular problem as there are virtually infinite permutations of pixel groups that are not keyways. For this keyway detection framework, precision was calculated to be 92.4% with a recall of 87.7% yielding an F1 score of 90.1%. This detection framework may also be applied to every brick interface in a full chanorama and the results are summarized in Figure 11.

---

<sup>c</sup>In the unlikely case that the HMT was a perfect fit for every or a large number of keyways, the optimization routine would determine a lower threshold of zero.

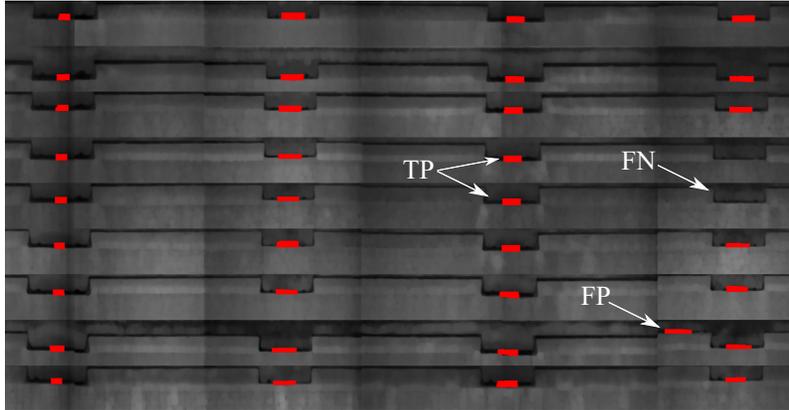


Fig. 11. Summary image showing all the interfaces of a single chanorama with all keyways detected using the HMT highlighted in red.

TABLE II  
Semantic segmentation keyway detection results.

		Actual	
		Positive	Negative
Prediction	Positive	TP = 347	FP = 1
	Negative	FN = 4	TN: NA

### III.B. Semantic Segmentation

We consider the exact same training and test dataset and performance metrics when evaluating the performance of the semantic segmentation keyway detection framework. Given in Table II is a summary of the performance for this approach. The only post processing rule required to improve performance is the suppression of detection objects in the binary detection mask with an area below a threshold which can be determined empirically. Using this approach precision was calculated to be 99.7% with a recall of 98.9% yielding an F1 score of 99.2%, all considerably better than achieved using the optimized HMT approach. This detection framework was also applied to every interface of the same chanorama as Figure 11 and this yields Figure 12. It is clear by comparing Table II and Table I that semantic segmentation has performed significantly better. Comparing Figure 11 and 12, we see that semantic segmentation detects the two keyways that the HMT missed. Accurately locating every keyway within the chanorama significantly reduces the complexity of analyzing the keyway positions to determine further information about the chanorama image, the placement of strips, etc.

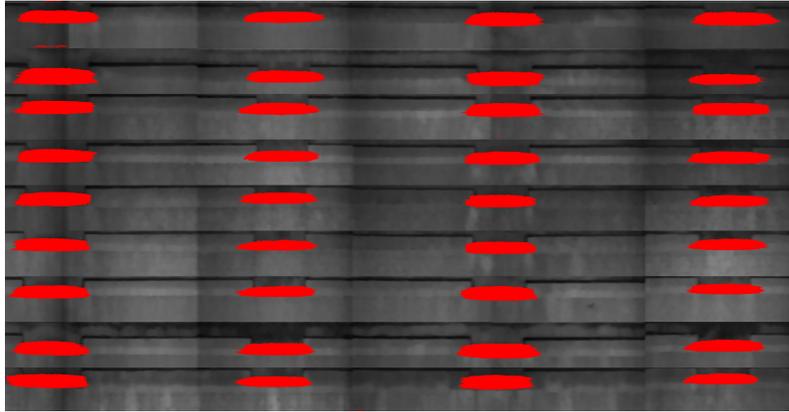


Fig. 12. Summary image showing all the interfaces of a single chanorama with all keyways detected using semantic segmentation highlighted in red.

### III.C. Keyway Dimensions

To calculate the performance of the algorithm to detect the sides of the keyways we consider the same test set used above but disregard any incomplete keyways. As each keyway has two sides this results in a test set of 598 keyway sides. Only exact detections of the keyway sides are considered. Out of the 598 keyways sides, 567 were detected exactly resulting in an accuracy of 94.8%. Of the missed detections, 27 occurred on one side of a keyway and only in two cases did the algorithm fail to detect either side of a keyway correctly. Some example results are shown in Figure 13. The example on the right of this figure is particularly challenging due to the seams between the strip that form a chanorama but the algorithm detects the correct edge regardless. Most chanoramas will have at least one column of keyways which exist on chanorama seams so it is important that the algorithm is robust. In a very small minority of cases, the interface strip results in a false detection such as the case as shown in Figure 14.

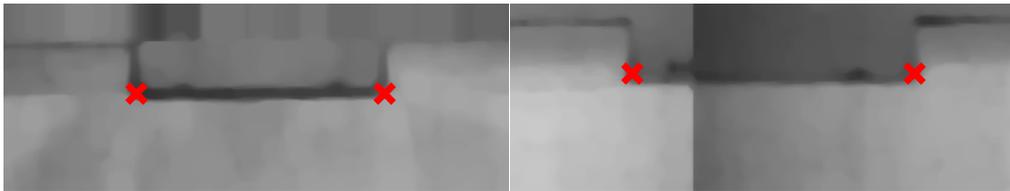


Fig. 13. Examples of side walls automatically detected.

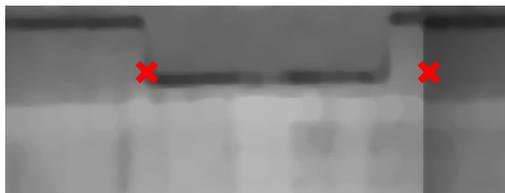


Fig. 14. Example of a case where the side wall is not correctly determined.

### III.D. Execution Time

To determine the time saved using the automated methods presented above, the execution time for each method was recorded while applying each framework to the entire test set. The Mathematical Morphology approach takes on average 0.267 seconds to detect every keyway in an interface strip (each interface strip contains 4 keyways) so approximately 0.067 seconds per keyway. The semantic segmentation approach takes on average 1.77 seconds to detect every keyway in an interface strip so approximately 0.443 seconds per keyway. A manual approach can take between one and two minutes to measure each keyway accurately if you take in to account the time required to open the very large chanorama file (each chanorama is approximately 180 Megapixels in size), locate the keyway of interest and make the measurement. The time taken for this approach is very dependent on the hardware resources available but even for high end hardware it could take many hours to measure every keyway in a large set of chanoramas recorded during a single inspection compared with a few minutes to apply the automated approach.

## IV. CONCLUSIONS

In this paper we have presented two unique approaches for automatically locating keyways in chanorama images and a technique for very accurately measuring the keyway dimensions using the output of the localization algorithms. Both approaches detect the keyways in a robust and reliable manner. Applying the keyway measurement algorithm after locating the keyways using either method allows for the accurate measurement of the dimensions of every keyway in a single chanorama in under 30 seconds. This activity would take considerably longer for a human operator to replicate manually and is a very laborious process for large datasets of chanoramas. This could result in a significant time saving when calculating scaling factors for sizing defects in chanoramas.

Each of the techniques has advantages and disadvantages. The advantages of the HMT

are the relative ease in the design of the detection framework and a similar technique could be readily designed for different problems with limited data. The disadvantages of this technique are its performance as it does not always detect every keyway which makes further analysis with keyway position difficult due to missing data. The main advantage of using the semantic segmentation approach is its excellent performance to detect nearly every keyway. The only examples where it fails to detect a keyway are where the interface strip was not cropped correctly when creating the dataset. The false positive rate is also very low. The disadvantages of using semantic segmentation are the redeployment of these pretrained models and detection frameworks for new problems and the lack of transparency on how it works which is a very important safety consideration in the nuclear industry. Both detection frameworks outlined in this paper locate the keyways but do not provide exact binary detection masks of the keyway locations. We demonstrated an accurate technique for locating the exact position of the sidewalls of the keyways using seed locations determined using either detection technique.

While we focused on applying the techniques outlined in this paper to only detect keyways regularly placed in panoramas images, both detection frameworks could readily be applied to detecting different types of physical features in different image types.

## **ACKNOWLEDGEMENT**

The authors would like to thank EDF Energy for providing the data used in this paper.

## REFERENCES

- [1] A. YOUNG, C. BERRY, G. M. WEST, and S. D. MCARTHUR, “A generalized model for fuel channel bore estimation in AGR cores,” *Nuclear Engineering and Design* (2019).
- [2] A. YOUNG, W. AYLWARD, P. MURRAY, G. WEST, and S. MCARTHUR, “Automatic anomaly detection in fuel grab load trace data using a knowledge-based system vs. multiple deep autoencoders,” (2019).
- [3] M. TRELINSKI, “Inspection of CANDU reactor pressure tubes using ultrasonics,” *17th World Conference on Nondestructive Testing*, 25–28 (2008).
- [4] M. TRELINSKI, “Ultrasonic testing methods and procedures for volumetric and surface inspection of CANDU reactor pressure tubes,” *European conference on Non-Destructive Testing* (2006).
- [5] M. PAJNIĆ, K. MARKULIN, A. MATOKOVIC, and H. FRANJIĆ, “Advanced approach of reactor pressure vessel in-service inspection,” *10th European Conference on Non-Destructive Testing* (2010).
- [6] S. GLASS, M. LEVESQUE, G. ENGELS, F. KLAHN, and D. FAIRBROTHER, “Underwater robotic tools for nuclear vessel and pipe examination,” *Proceedings ANS 8th International Topical Meeting on Robotics and Remote Systems*, 25–29 (1999).
- [7] B.-H. CHO, S.-H. BYUN, C.-H. SHIN, J.-B. YANG, S.-I. SONG, and J.-M. OH, “KeproVt: underwater robotic system for visual inspection of nuclear reactor internals,” *Nuclear engineering and design*, **231**, 3, 327 (2004).
- [8] S.-U. LEE, Y.-S. CHOI, K.-M. JEONG, and S. JUNG, “Development of a Tele-operated Underwater Robotic System for maintaining a light-water type power reactor,” *2006 SICE-ICASE International Joint Conference*, 3017–3021, IEEE (2006).
- [9] C.-H. CHOI, S. H. KIM, C.-H. KIM, Y.-C. SEO, H. C. SHIN, S.-U. LEE, K.-M. JUNG, S.-H. JUNG, and Y. S. CHOI, “The Robots for Nuclear Power Plants,” *Conference on Korea Nuclear Society 2005* (2005).

- [10] A. COLE-BAKER and J. REED, *Measurement of AGR graphite fuel brick shrinkage and channel distortion*, vol. 309, 201–208, Royal Society of Chemistry.
- [11] P. MURRAY, G. WEST, S. MARSHALL, and S. MCARTHUR, “Automated in-core image generation from video to aid visual inspection of nuclear power plant cores,” *Nuclear Engineering and Design*, **300**, 57 (2016).
- [12] S. J. SCHMUGGE, N. R. NGUYEN, C. THAO, J. LINDBERG, R. GRIZZI, C. JOFFE, and M. C. SHIN, “Automatic detection of cracks during power plant inspection,” *Proceedings of the 2014 3rd International Conference on Applied Robotics for the Power Industry*, 1–5, IEEE (2014).
- [13] M. R. JAHANSHAHI, S. F. MASRI, C. W. PADGETT, and G. S. SUKHATME, “An innovative methodology for detection and quantification of cracks through incorporation of depth perception,” *Machine vision and applications*, **24**, 2, 227 (2013).
- [14] S. J. SCHMUGGE, L. RICE, N. R. NGUYEN, J. LINDBERG, R. GRIZZI, C. JOFFE, and M. C. SHIN, “Detection of cracks in nuclear power plant using spatial-temporal grouping of local patches,” *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1–7, IEEE (2016).
- [15] H. OLIVEIRA and P. L. CORREIA, “CrackIT - An image processing toolbox for crack detection and characterization,” *2014 IEEE international conference on image processing (ICIP)*, 798–802, IEEE (2014).
- [16] S. J. SCHMUGGE, L. RICE, J. LINDBERG, R. GRIZZI, C. JOFFE, and M. C. SHIN, “Crack segmentation by leveraging multiple frames of varying illumination,” *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1045–1053, IEEE (2017).
- [17] F.-C. CHEN and M. R. JAHANSHAHI, “NB-CNN: deep learning-based crack detection using convolutional neural network and Naive Bayes data fusion,” *IEEE Transactions on Industrial Electronics*, **65**, 5, 4392 (2017).
- [18] F.-C. CHEN, M. R. JAHANSHAHI, R.-T. WU, and C. JOFFE, “A texture-based video processing methodology using Bayesian data fusion for autonomous crack detection on metallic surfaces,” *Computer-Aided Civil and Infrastructure Engineering*, **32**, 4, 271 (2017).

- [19] R. SUMMAN, N. MATHUR, G. DOBIE, G. WEST, S. MARSHALL, C. MACLEOD, C. MINEO, and G. PIERCE, “A new probe concept for internal pipework inspection,” (2017).
- [20] M. DEVEREUX, P. MURRAY, G. WEST, S. BUCKLEY-MELLOR, G. COCKS, C. LYNCH, and A. FLETCHER, “Automated analysis of AGR fuel channel inspection videos,” *6th EDF Energy Nuclear Graphite Conference* (2018).
- [21] Z. ZHANG, “A flexible new technique for camera calibration,” *IEEE Transactions on pattern analysis and machine intelligence*, **22** (2000).
- [22] J. SERRA, *Image analysis and mathematical morphology, v. 1*, Academic press (1982).
- [23] N. R. HARVEY, R. PORTER, and J. THEILER, “Ship detection in satellite imagery using rank-order grayscale hit-or-miss transforms,” *Visual Information Processing XIX*, vol. 7701, 770102, International Society for Optics and Photonics (2010).
- [24] B. PERRET, S. LEFÈVRE, and C. COLLET, “A robust hit-or-miss transform for template matching applied to very noisy astronomical images,” *Pattern Recognition*, **42**, 11, 2470 (2009).
- [25] N. R. HARVEY and S. MARSHALL, “Using genetic algorithms in the design of morphological filters,” *Mathematical Morphology and Its Applications to Image Processing*, 53–59, Springer.
- [26] N. R. HARVEY and S. MARSHALL, “The use of genetic algorithms in morphological filter design,” *Signal Processing: Image Communication*, **8**, 1, 55 (1996).
- [27] V. BADRINARAYANAN, A. KENDALL, and R. CIPOLLA, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *arXiv preprint arXiv:1511.00561* (2015).
- [28] S. AKÇAY, M. E. KUNDEGORSKI, M. DEVEREUX, and T. P. BRECKON, “Transfer learning using convolutional neural networks for object classification within x-ray baggage security imagery,” *2016 IEEE International Conference on Image Processing (ICIP)*, 1057–1061, IEEE (2016).

- [29] E. ROSTEN and T. DRUMMOND, “Fusing points and lines for high performance tracking,” *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, vol. 2, 1508–1515, Ieee (2005).
- [30] J. SHI ET AL., “Good features to track,” *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, 593–600, IEEE (1994).
- [31] C. G. HARRIS, M. STEPHENS ET AL., “A combined corner and edge detector.” *Alvey vision conference*, vol. 15, 10–5244, Citeseer (1988).
- [32] R. C. GONZALEZ and R. E. WOODS, *Digital Image Processing (3rd Edition)*, Prentice-Hall, Inc. (2006).
- [33] P. SOILLE, *Advances in the Analysis of Topographic Features on Discrete Images*, 175–186, Springer Berlin Heidelberg, Berlin, Heidelberg.
- [34] P. SOILLE, *Morphological Image Analysis: Principles and Applications*, Springer-Verlag New York, Inc. (2003).
- [35] D. BEASLEY, D. R. BULL, and R. R. MARTIN, “An overview of genetic algorithms: Part 1, fundamentals,” *University computing*, **15**, 2, 56 (1993).
- [36] K. DEB, “An introduction to genetic algorithms,” *Sadhana*, **24**, 4-5, 293 (1999).
- [37] K. DEEP, K. P. SINGH, M. L. KANSAL, and C. MOHAN, “A real coded genetic algorithm for solving integer and mixed integer optimization problems,” *Applied Mathematics and Computation*, **212**, 2, 505 (2009).
- [38] G. J. BROSTOW, J. FAUQUEUR, and R. CIPOLLA, “Semantic object classes in video: A high-definition ground truth database,” *Pattern Recognition Letters*, **30**, 2, 88 (2009).
- [39] A. KRIZHEVSKY, I. SUTSKEVER, and G. E. HINTON, “Imagenet classification with deep convolutional neural networks,” *Advances in neural information processing systems*, 1097–1105 (2012).
- [40] H. BAY, T. TUYTELAARS, and L. VAN GOOL, “Surf: Speeded up robust features,” *Computer vision ECCV 2006*, 404–417 (2006).

- [41] I. GOODFELLOW, Y. BENGIO, A. COURVILLE, and Y. BENGIO, *Deep learning*, vol. 1, MIT press Cambridge (2016).