

Using Generative Adversarial Networks to Improve the Efficiency of Crack Detection in Nuclear Reactor Inspection Data

E. Branikas¹, P. Murray¹, and G. West¹

¹University of Strathclyde, Glasgow, United Kingdom

efstathios.branikas@strath.ac.uk, paul.murray@strath.ac.uk, Graeme.west@strath.ac.uk

Abstract

In the nuclear sector, condition monitoring & inspection activities allow operators to assess and understand the health and condition of their plant. In the UK, where power stations typically belong to the Advanced Gas-cooled Reactor (AGR) family, inspections of the reactor cores are performed using various sensors including cameras which allow remote visual inspection (RVI) of key components. Specifically, RVI of the reactor core allows operators to view the inner surface of individual fuel channels selected for inspection. In this process, key reactor components are assessed, and any potential defects or anomalies can be detected and classified. Analysis of the resulting video footage is typically conducted manually, but recently, new approaches capable of identifying defects automatically, have been proposed to provide new decision support tools for analysing the data. However, current state-of-the-art techniques presented in literature have their limitations, particularly in terms of their ability to accurately detect and describe features such as cracks. This paper therefore proposes a new approach which addresses this challenge and is demonstrated through a case study using representative data gathered from AGR inspections. To achieve our goal, we implement a neural network that belongs to the family of Generative Adversarial Networks (GANs). Our approach provides a novel post-processing step which improves the pixel-level prediction of an existing image segmentation network designed to detect and segment cracks. We also provide a framework that can be incorporated into the output of other algorithms which produce binary crack predictions at their output. Our approach is demonstrated through application to Remote Visual Inspections of graphite core bricks in Advanced Gas-cooled Reactors and has the potential to be applicable to other inspection activities in the nuclear sector and other industries.

1. Introduction

In the case of Advanced Gas-Cooled Reactor (AGR) power stations, graphite bricks are used both to house the control rods and the fuel of the reactor, and to allow the flow of CO₂ gas that acts as coolant. The core is comprised of hollow cylindrical graphite bricks. The interlocking graphite bricks have a structural function to maintain the core's geometry under applied loads to ensure free movement of fuel, control rods and coolant. The graphite also plays the role of the moderator which propagates the fission reaction. Inspection of selected fuel channels in the reactor core is conducted during planned, periodic outages. Various sensors are used during inspection including a video camera which allows Remote Visual Inspection (RVI). During RVI, an inspection tool known as NICIE (New In-core Inspection Equipment), equipped with a camera is inserted in the fuel channels to be inspected and data is gathered from the entire channel bore at 6 different overlapping orientations. The inspection footage is recorded and further analysed by experts who assess the data and evaluate whether the station is safe to return to operation. ASIST (Automated Software Image Stitching Tool) software [1] [2] has been developed to produce stitched montages

of the whole fuel channel automatically. This process stitches different orientation footage of the raw video and combines it seamlessly to produce 360° views of the channel, which are called “chanoramas” standing for channel-panoramas.

Chanoramas have been used extensively for manual assessment, health inspection, and defect quantification of the fuel channels. However, this process is conducted by human experts and is manually intensive, often prone to subjectivity and the results of the analysis and crack detection are difficult to reproduce. Recently, new automated approaches have been developed for analysing AGR inspection data [3] [4]. However, many challenges posed by the nature of defects and crack detection do not allow state-of-the-art models to perform optimally [3]. Cracks and defects only constitute a small fraction of the overall image, creating a great data imbalance that is challenging to overcome. The AGR case introduces additional challenges including very different visual textures, varying illumination conditions and differences in contrast between cracks and background. Figure 1 shows examples of small sections of the inner surface of AGR fuel channels depicting cracks. These images are extracted from chanoramas and illustrate some of the challenges mentioned above.

Nuclear power stations typically require more frequent inspections as their age increases, and specifically in the UK, where AGRs are approaching the end of their expected lifetime. Regular inspections are therefore taking place to discover, evaluate, and monitor potential structural defects in various elements of the station and to support station operators’ decision making around the continued safe operation of the plant. Typically, the analysis of the inspection data is conducted manually, and as it is a labour-intensive process, research is exploring new automated methods to provide decision support tools for analysing the data acquired during inspection [3]. In this paper, we present a new post-processing model that improves the detection accuracy of automated crack detection methods applied to inspection footage from the inner surface of AGR components. Although we showcase a specific inspection process, our method can theoretically be applied to similar inspection processes where the output is a pixel-level detection or segmentation map of cracks and other defects.

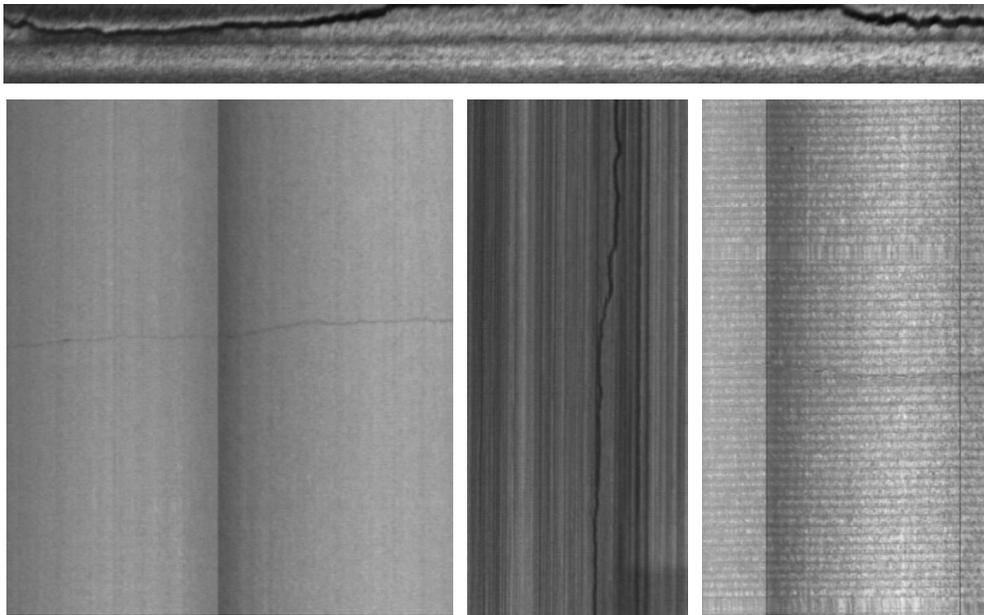


Figure 1: AGR examples of cracks on graphite bricks.

Despite the challenges described, new approaches are being developed for automated crack and defect detection in nuclear power inspections. In [4], Devereux et al. developed a method for detection and quantification of cracks on panoramas. The authors provide patch-level detection and further localize the area of interest by iteratively providing the model with smaller surfaces as an input. A similar approach [5] utilizes a Convolutional Neural Network (CNN) to detect defects and cracks on nuclear inspection footage and through a spatiotemporal registration. In [5], multi-frame information of a given surface is exploited with a stochastic fusion operation, to acquire higher detection accuracy on a patch level of images. The authors of [6], similarly perform patch-level classification of cracks on nuclear inspection data from multi-frame information and subsequently use thresholds to obtain the patch prediction. Pixel-level detection of cracks in nuclear inspection footage has also been explored. The authors of [3] deploy U-Net [7] a well-known CNN for the task of pixel-level crack prediction, also known as image segmentation. This model is tested on panoramas originating from video footage acquired from RVI. Our new model, which is proposed in this paper, is tested on the same dataset as in [3] and applied on top of the approach used in that work. Neural networks for pixel-level prediction are able to detect complex visual patterns that are available on the training examples but can fail to perform equally well on previously unseen cases. To address this, our model is used as a post-processing step to enhance the detection capabilities of the U-Net model, using a Generative Adversarial Network (GAN).

GAN models have attracted a lot of attention since their early conception [8]. Originally, they were created for generating artificial data, and their superior generative capabilities have led in development of GAN models for various purposes, ranging from image stylization [9], to super-resolution [10] and image translation [11] [12]. Adversarial models have been introduced in fault detection and monitoring as well [13] [14]. In [13] a generative model is deployed to improve the classification accuracy of still defects. The authors of [14] deploy CycleGAN [12], a famous model used for image translation to perform pixel-level detection of cracks without respective binary labels.

In this paper, we propose a novel post-processing method for improving the pixel-level detection of cracks in inspection data of AGR power stations. To achieve this, we deploy and train Pix2Pix [11], a GAN model used for paired image-to-image translation purposes. The model is trained with small patches extracted from panoramas, and the respective imperfect predictions of the U-Net model used in [3]. Our post-processing GAN is trained to map our original imperfect crack segmentations more accurately to the ground-truth images.

In summary, the main contributions of this work can be summarized as creating a novel post-processing method for improving the crack detection capabilities of an established segmentation model. We demonstrate the usefulness of our approach using AGR data although method has potential to be applied to inspections carried out in different components of nuclear power stations and in other industries. The rest of the paper can be summarised as follows: In Section 2 we provide all the required background for Generative Adversarial Networks and the proposed approach; in Section 3 we present both qualitative and quantitative results of the experiments conducted, and we also discuss the analyse them, before drawing our conclusions in Section 4.

2. Defect Detection Refinement Framework

2.1 Generative Adversarial Networks

GAN models are based on an adversarial learning scheme. In brief, there are two competing neural networks, the Generator (G) and the Discriminator (D), playing a minimax game against each other. The Generator’s goal is to create an image, that is as close to the real distribution of data as possible, while the Discriminator, being a classifier, distinguishes whether the input it accepts is a real piece of data or a generated example. In this way, the Generator adapts and updates its variables based on the feedback of the Discriminator. Equally, the Discriminator gradually improves its ability to discriminate between real and artificial data, through the optimization process. These concepts are illustrated in Figure 2.

It is critical that the two networks have similar learning capabilities and are learning at an equal rate, otherwise instability will hinder the model; if the Discriminator surpasses the Generator, it will always predict the correct class and not allow the Generator to adapt in a way that will *trick* the Discriminator, thus creating more realistic outputs. On the contrary, if the Discriminator is tricked from the Generator often, it will not provide useful feedback that allows the Generator to improve further. The whole process is completed when an equilibrium state is reached, and the Discriminator provides equal probabilities for a generated image, to be real or fake.

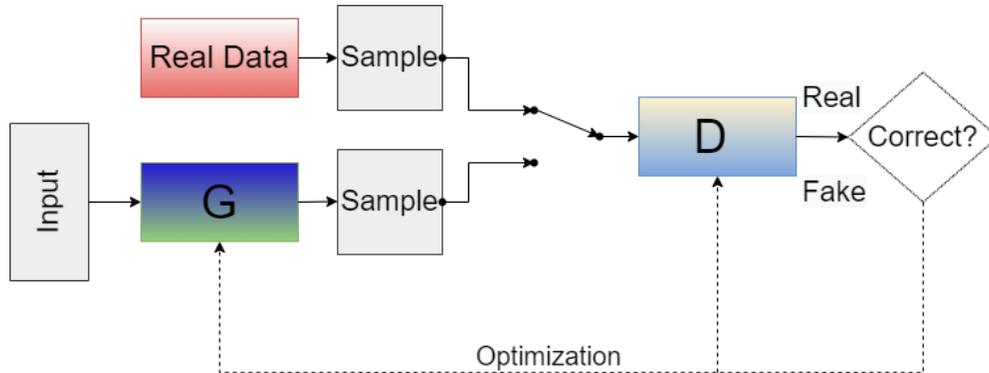


Figure 2: The fundamental framework of a Generative Adversarial Network. The discriminator is trained in an alternating fashion with real and generated examples.

The two models are trained together in an adversarial manner, by optimizing a common objective function. Given an input z that can either be a variable from a random distribution or an image in the case of image-to-image translation problems, and x an example from the real distribution of data, a common adversarial loss is defined as:

$$L_{adv}(G, D) = \mathbb{E}_{x \sim p_{data}(x)} [\log(D(x))] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))] \quad (1)$$

Equation (1) presents the adversarial loss which is a binary cross-entropy function commonly used in GAN models [8]. The optimization of the two networks is achieved by updating the equation as $\min_G \max_D L_{adv}(G, D)$. In this way, the Generator is adjusted through the training process so that

$D(G(z))$ is maximized and respectively the Discriminator is adjusted so that $D(x)$ is maximized and $D(G(z))$ is minimized.

2.2 Pix2Pix Model

To create a model that improves the segmentation results of crack detection on nuclear inspection data, we select Pix2Pix [11], a well-known GAN model that is deployed for image-to-image translation problems. Image-to-image translation refers to adjusting an image from one domain to another, in a way that there is correspondence between the two domains. Pix2Pix requires paired examples between the original and the target domain. In our method, this refers to an imperfect crack prediction binary image that derives from a segmentation method, and its counterpart, which is the humanly annotated ground-truth i.e., the *correct* segmentation according to a human expert. Given these pairs, the Generator is trained to map a sub-optimal prediction the human annotation.

The Generator of Pix2Pix for our application is U-Net [7], and has the same architecture as in [3]. Importantly, the Generator is trained on different inputs to the initial segmentation network [3] and, in this post-processing approach it serves the purpose of crack detection *refinement*. U-Net is an auto-encoding neural network that is used widely for a variety of applications. It consists of two parts, the encoder, that compresses the input into a new, feature-rich representation, and the decoder, that expands the resolution of the encoder’s output, and gradually reduces the number of feature channels. Skip connections between the respective layers of the encoder and the decoder aid in refining the flow of information and maintaining low-level information. Figure 3 depicts the overall architecture and layers of the U-Net model.

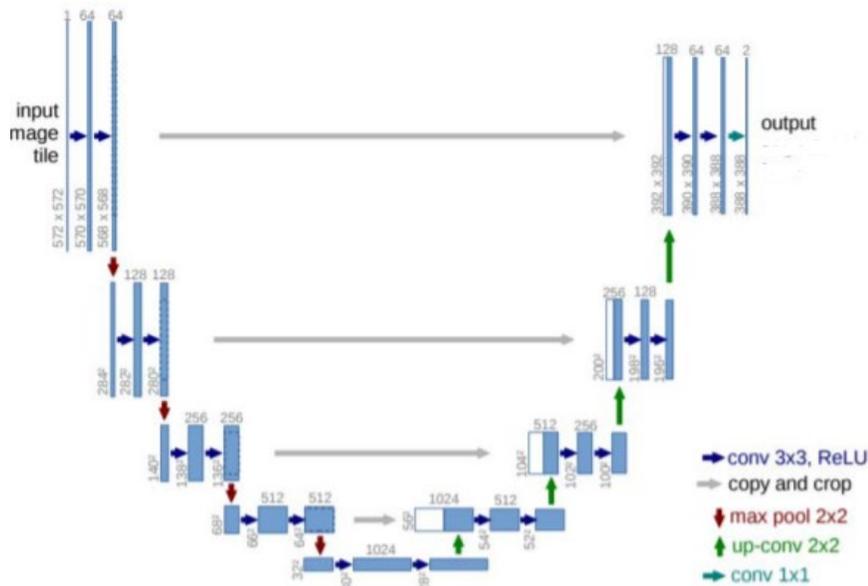


Figure 3: U-Net architecture. [7]

The Discriminator is of the same structure as the one that is used originally in Pix2Pix, a patchGAN classifier [11]. This network classifies an input in patches, and is applied convolutionally to an image, providing a matrix as an output. Every value of this matrix denotes the discriminator’s prediction of the respective patch being real or fake. The architecture is a stack of convolutions and batch normalization operations, followed by a non-linear ReLU function, which is the basic

block of our Discriminator. First, a 7 by 7 convolution block is applied, followed by 5 blocks of smaller kernels (3 by 3 in dimensions) and finally a single convolution operation with a 1 by 1 kernel is applied to ensure the correct number of channels is acquired at the output of the Discriminator.

2.3 Crack Detection Refinement

Based on the Pix2Pix model, the framework we propose is deployed for the crack refinement task without any further adjustments on the Pix2Pix model. The Generator is fed with patches of sub-optimal crack predictions produced from the model developed in [3], as can be seen in Figure 4, which serve as the *fake* data in the GAN setting. The Generator tries to map these images to the *real* distribution which is the manually annotated ground-truth images. Following this, the discriminator’s input is the output of the Generator. This output is concatenated with a condition which, in this case, is the original image of the inspection to which the *fake* predictions from our segmentation network correspond. In this way, the discriminator is trained by making use of both the original grayscale image and the initial segmentation map generated by U-Net for this image. Equally, when the Discriminator is optimized on real data, its input is the original grayscale image and the ground-truth. In our experiments, we also explore providing the Generator with the grayscale image information, but the semantic gap between inspection images and binary labels is found to be too great thus preventing the generator from producing a refined crack prediction.

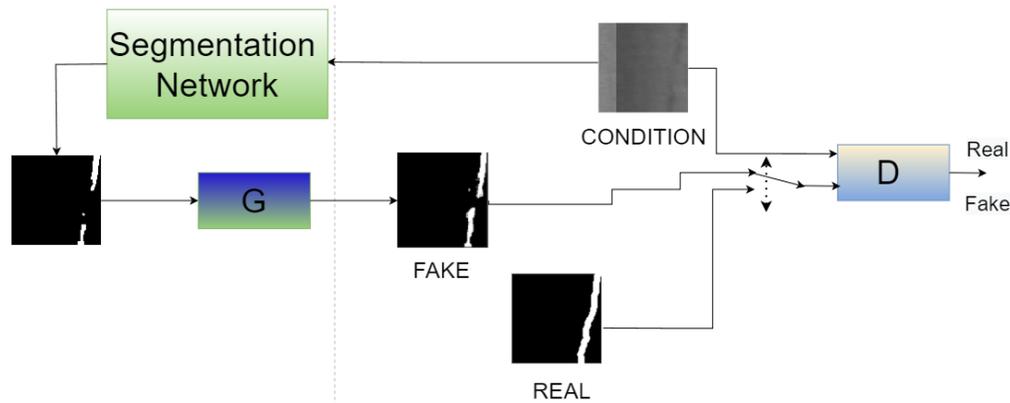


Figure 4: Training setting of the refinement process. The Generator’s input is only a faulty prediction, while the Discriminator’s input also includes the respective inspection image.

After the training process is complete, the model is deployed by feeding the Generator with only faulty predictions from an initial segmentation. Thus, the generator effectively operates as a post-processing step which aims to improve the output of our initial segmentation. Our approach is agnostic to the original image during testing, as only the Generator is required. This offers the potential for reusability on the segmentation results of different inspection data as well.

3. Experimental Results

To assess our framework, we evaluate the GAN model on the crack detection results of a U-Net model, as defined in [3]. In that work, a dataset of 109 images of varying sizes originating from panoramas was assembled. 11 of these images were used for testing and the remaining for training and evaluating the network. To provide a fair evaluation of our approach, we use the same number of images for training, evaluating, and testing our model, as in in [3], where a U-Net model was used for image segmentation. For training and evaluating our proposed GAN model, patches of 96 by 96 pixels in dimensions were extracted, only when crack predictions were present in the patches, with an overlap of 12 pixels in a sliding window manner. This resulted in 47000 patches for which the: 1) U-Net predictions; 2) ground truth labels; 3) original image surfaces, were extracted to create data triplets. Each data triplet corresponds to the same spatial coordinates.

After training the Pix2Pix model, as described in Section 2, an assessment of the refined outputs of the generator is provided, for the testing images. The per-pixel evaluation is based on the same metrics that were used for evaluating the original crack segmentation U-Net model:

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive} \quad (2)$$

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negative} \quad (3)$$

$$F1\ Score = \frac{2 * Precision * recall}{precision + recall} \quad (4)$$

Following the same approach as in [3] and [15], in order not to penalize for slightly thicker crack prediction of the models, we allow a radius of 5 pixels of margin around cracks, for which false positives are considered true positives.

Table 1: Test evaluation metrics for the original predictive model (U-Net) and the proposed refinement Pix2Pix framework.

Column 1 Title	Direct pixel comparison			Comparison with 5-pixel compensation		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score
U-Net predictions	0.872	0.787	0.827	0.988	0.807	0.889
Pix2Pix Predictions	0.865	0.804	0.833	0.969	0.827	0.893

The results shown in Table 1 suggest our proposed model improves the crack detection results of the U-Net segmentation model, providing superior performance. Of the 3 metrics, recall is the one that refers to the ratio of the detected crack pixels. For applications like nuclear inspection, this is the most critical factor, as it reflects to the model’s ability to successfully detect the object of interest, which, in our application, is underrepresented in the dataset. The goal of this research is to increase the detectability of cracks and other defects. As such, the increase in the recall metric that is observed suggests the proposed model indeed achieves this goal. As there is a trade-off between this metric and precision, as expected, the precision of the refinement approach is slightly lower than in the original segmentation. Our proposed post-processing technique allows more crack pixels to be detected accurately, but this comes at the expense of a few false positives being

introduced. Ultimately, the overall F1 score, which provides the overall quality of the detection, is higher for our proposed new GAN model. This result provides confidence that our new post-processing approach increases the overall segmentation quality.

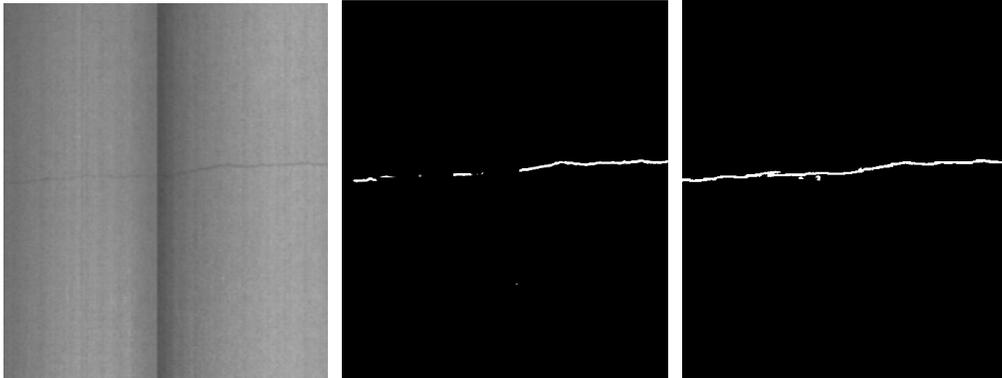


Figure 5: Example of a challenging case of a defect and the original U-Net prediction (middle) against the improved version of the proposed model (right).

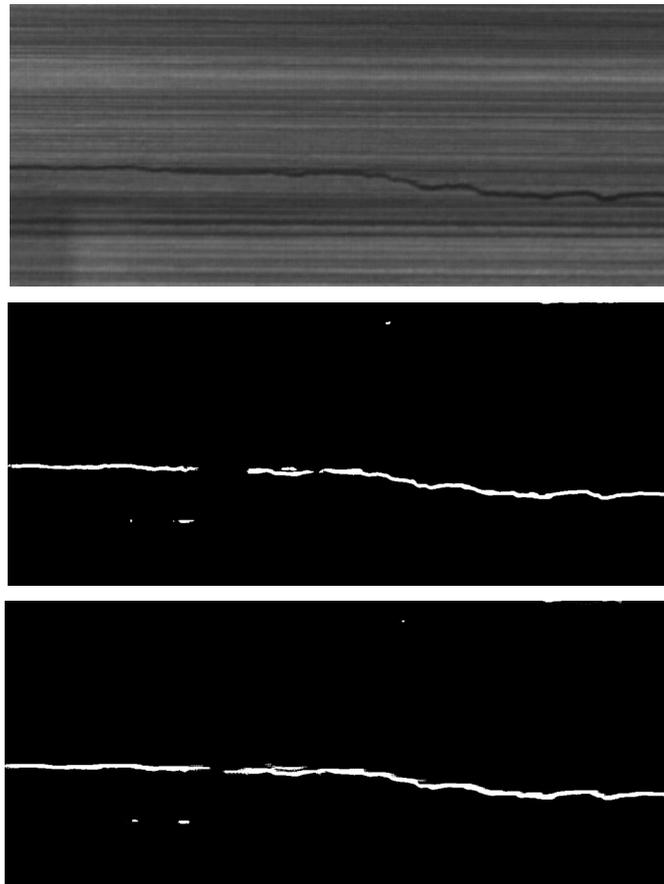


Figure 6: Example of a challenging case of a defect and the original U-Net prediction (middle) against the improved version of the proposed model (bottom). The proposed model does not enhance the noise on the U-Net model, but detects the false negatives and adjusts the crack prediction accordingly.

Figure 5 and Figure 6 show some images which allows a qualitative comparison between our framework and the original predictions. It is evident that the GAN-derived binary masks succeed in partially or fully restoring the missing parts of the undetected defects and in parallel do not introduce additional noise that might appear as isolated pixels. These findings confirm the quantitative results presented in Table 1. This further improves our confidence that the proposed algorithm has the potential to be used as a plug-and-play post-processing step in addition to an initial image segmentation module.

4. Conclusions

In this paper, we have presented a new GAN-based model that can be used as a post-processing step for refining pixel-level predictions and segmentations of cracks in RVI data of real AGR fuel channels. We have implemented the Pix2Pix model in such a way that after training, the model can provide refined and improved binary segmentation maps when the input comes from the output of an existing crack detection model. By drawing a comparison with the previous method [3], we demonstrate that our GAN based post-processing step improves the crack detection performance without inserting substantial false positive error. Although it could be argued that our approach is not making good use of the available information e.g., it does not incorporate information of the original image during deployment, we believe that this helps our model generalise better. As a result, this opens up the possibility of directly applying this trained neural network on inspections from different nuclear station components or for similar inspection challenges in other sectors outside of nuclear. In the future, modification of the Generator architecture can be considered, to allow making use of the original inspection visual information as well.

5. Acknowledgments

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

6. References

- [1] 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*, vol. 300, p. 57–66, 2016.
- [2] G. West, P. Murray, S. Marshall and S. McArthur, "Improved visual inspection of advanced gas-cooled reactor fuel channels," *International Journal of Prognostics and Health Management*, vol. 6, 2015.
- [3] E. Branikas, P. Murray and G. West, "An innovative crack detection algorithm to support automated inspection of nuclear reactor cores," in *12th Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (NPIC&HMIT 2021)*, 2021.
- [4] M. G. Devereux, P. Murray and G. M. West, "A new approach for crack detection and sizing in nuclear reactor cores," *Nuclear Engineering and Design*, vol. 359, p. 110464,

2020.

- [5] F.-C. Chen and M. R. Jahanshahi, "NB-CNN: Deep learning-based crack detection using convolutional neural network and Naïve Bayes data fusion," *IEEE Transactions on Industrial Electronics*, vol. 65, p. 4392–4400, 2017.
- [6] S. J. Schmutge, 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," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [7] O. Ronneberger, P. Fischer and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, 2015.
- [8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [9] T. Karras, S. Laine and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019.
- [10] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang and others, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [11] P. Isola, J.-Y. Zhu, T. Zhou and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [12] J.-Y. Zhu, T. Park, P. Isola and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017.
- [13] H. Di, X. Ke, Z. Peng and Z. Dongdong, "Surface defect classification of steels with a new semi-supervised learning method," *Optics and Lasers in Engineering*, vol. 117, p. 40–48, 2019.
- [14] K. Zhang, Y. Zhang and H.-D. Cheng, "Crackgan: Pavement crack detection using partially accurate ground truths based on generative adversarial learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, p. 1306–1319, 2020.
- [15] M. D. Jenkins, T. A. Carr, M. I. Iglesias, T. Buggy and G. Morison, "A deep convolutional neural network for semantic pixel-wise segmentation of road and pavement surface cracks," in *2018 26th European Signal Processing Conference (EUSIPCO)*, 2018.