Image Stitching and Illumination Correction for Metallic Multi-Purpose Canister Visual Inspection Data

Michael Devereux, Paul Murray and Graeme West

University of Strathclyde
204 George Street, Glasgow, G1 1XW, U.K.
michael.devereux@strath.ac.uk; paul.murray@strath.ac.uk; graeme@strath.ac.uk

INTRODUCTION

Visual inspection is an important activity common in the assessment of various assets in the nuclear industry. Visual inspection has applications ranging from determining the health of the nuclear reactor core itself [1, 2], determining the health of vital components of the nuclear power plant such as superheaters [3], assessing the conditions of welds in thick-walled pressure vessels common to Boiling Water Reactor (BWR) and Pressurized Water Reactors (PWR) [4] and many other applications. Each of these activities typically involves gathering a large amount of raw data either in the form of video or still images. At any given instance of this data it is only possible to resolve localised information therefore suitable processing is required to generate a global view of the asset.

In this paper we focus on visual inspections of multi-purpose canisters (MPC) used for the storage of waste from PWR operation. We present an approach for addressing two problems encountered during the analysis of raw inspection data, namely stitching multiple images together to form a single montage and correcting for non-uniform illumination from the inspection rig.

For the inspection of the MPCs, a specialist camera rig captures thousands of high resolution images (24MP) which results in each inspection generating a very large dataset. This problem is similar to [2, 5] where they develop specialised software for the generation of inspection montages spanning the entire inner surface of fuel channels from routine inspection video. The main difference with this approach is that we are working with high resolution photographs instead of inspection videos. To aid in the inspection process and assessment of the large amount of data, a series of novel algorithms suitable for stitching this particular type of data have been developed. These algorithms take each of the photographs, calculates the relative position of each of the image tiles and stitches all the images together with the appropriate amount of overlap to generate a large composite image which spans the entire outer surface of the MPC.

The second problem this paper addresses is the non-uniform illumination present on each of the inspection images. Non-uniform illumination is a very common problem in image processing and impacts a diverse set of fields ranging from fingerprint analysis[6], transmission electron microscopy[7] to document binarization[8]. Images recorded of a metallic curved surface are prone to reflection and lighting issues therefore obtaining uniform and consistent lighting for all the inspection images is not possible without significant post-processing. We describe an innovative approach to learning a generalised lighting model for the imaging system. This lighting model can than be applied to all inspection images to correct the non-uniform illumination. This model or illumination profile is an important step in providing surface inspection images which are more uniform and consistent.

While this paper focuses on stitching MPC inspection images and correcting the illumination, the techniques described could easily be adapted to stitch visual inspection data from inspections of other types of container both in the nuclear industry and and outwith where routine inspections are performed.

Inspection data

The inspection rig consists of a DSLR[1] camera with a prime lens and a 2D translation stage which moves the camera to regular positions and takes a photograph at each position in a manner that ensures overlap between photographs. This automated approach is used to ensure the image gathering process remains constant from inspection to inspection and that complete coverage of the surface is achieved. The DSLR camera is equipped with a flash unit to provide lighting and photographs are taken at each position of interest on the surface of the MPC. Sample inspection images are shown in Fig. 1. The camera parameters are kept consistent for all the photographs to maintain consistency throughout the final stitched montage but due to the complex interactions of light with the metallic surface, features on the surface of the MPC can appear different when illuminated or viewed from different angles. This is further exacerbated by non-uniformities of the lighting profile across the curved surface of the MPC. This effect is far more pronounced on some images where the metal surface is relatively shiny compared with other regions of the MPC.

Fig. 1: Sample inspection images.

IMAGE STITCHING

In this section we outline the methods used to stitch the inspection data described in the previous section. We have

[1] Digital single-lens reflex camera

This is a peer-reviewed, accepted author manuscript of the following paper: Devereux, M., Murray, P., & West, G. (Accepted/In press). Image stitching and illumination correction for metallic multi-purpose canister visual inspection data. In Decommissioning Environmental Science and Remote Technology 2021 American Nuclear Society.

investigated the application of traditional stitching algorithms to this data but due to the problems described above and the lack of salient features in the majority of the footage, typical feature detection and image stitching algorithms do not work. However, it is still possible to stitch the images using positional data encoded in the files. The coordinates of each photograph is recorded in the file name of each image file and this can be automatically extracted as the first step in the stitching process. This positional data occasionally contains small errors due to the output of the translation stage stepper motor having limited accuracy so the algorithms have to be robust to these errors. The algorithms also need to be robust to missing data as occasionally the flash does not fire at a particular position or the quality of the image gathered was not deemed to be good enough[2].

The most robust approach to stitching the data together is the creation of a matrix of filenames where the location in the matrix represents where the image will be stored in the final stitched image. Generating this matrix is a nontrivial task due to the small errors in positional data in the file names and missing inspection images. The first stage in generating this filename matrix is generating a raw coordinate matrix which represents a regular 2D spatial grid. The x and y values to be used to generate this coordinate matrix may be calculated using density based spectral clustering [9] on every element of the x coordinate array and the same for the y coordinate array extracted from every filename in any particular dataset. This matrix is than used to generate the filename matrix by comparing every entry in the coordinate matrix with the positional data for every single image in the dataset. If the Euclidean distance between the image file and the coordinate values for any given entry in the coordinate matrix are under a certain predefined threshold, this filename will be assigned to the corresponding position in the filename matrix. The final step in the stitching process is to use the filename matrix to generate an inspection montage. This is achieved by first stitching all the images in every row together and than stitching the resulting columns together to give the final montage.

ILLUMINATION CORRECTION

In this section we describe an approach for correcting non-uniform illumination. Correcting any non-uniformities in the illumination will enhance the perceived visual quality of the inspection images and MPC surface maps and it will also allow more robust image analysis and particularly image comparisons and quantification of change. The first step in correcting these problems is to estimate a generalised model for the illumination provided by the imaging system. This may be achieved by sampling a large number of images (n), summing them all together and calculating the mean image. The generated mean image will contain both the illumination profile of the system but also noise and possibly other high frequency information which do not represent the lighting profile. To extract just the illumination profile, we use the arbitrary dimension smoothing algorithms described in [10, 11].

[2]When a photograph fails to be recorded at any position or the quality is not deemed sufficient, the camera is manually returned to this position and an inspection image is manually recorded

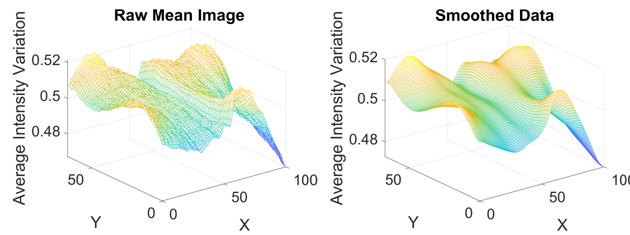

(a) Surface plot of raw mean image.

(b) Surface plot of smoothed mean image.

Fig. 2: Illumination surface plots before and after smoothing.

These algorithms take the mean image and remove the high frequency components and noise leaving just the low frequency components which represent the variation in lighting. Given in Fig. 2a is a surface plot of the mean image where high frequency characteristics are clearly present. Figure 2b is a surface plot of where 3 dimensional smoothing has been applied to the mean image. This resulting plot is the illumination profile which we use to correct all images in the dataset.

To test for generalisation the full data set containing N images should by subsampled multiple times with each sample containing n images. The resulting smoothed illumination profile should remain consistent for different subsets of each inspection. It is also possible to check if the lighting remains consistent between different inspections where are performed on different years. Ideally the illumination profile should remain consistent for different subsets of data from each inspection and consistent from inspection to inspection. This is demonstrated to be the case by comparing the plots in Fig. 3a and Fig. 3b. In this series of plots we see that the surface representing the illumination profile remains very consistent from inspection to inspection[3] but also when considering different subsets of the data from each inspection. These plots represent just the variation in lighting as the base level lighting has been subtracted from each plot for clarity. To correct the illumination in any inspection image, a mean illumination profile is calculated which represents all inspection data to date. This illumination profile can than be subtracted from each inspection image to give the light corrected image.

RESULTS

In this section we present the results of the application of the two approaches described so far to real data recorded during the inspection of the MPC.

Stitching

Given in Fig. 6 is the final stitched montage generated by the stitching algorithms. The black rectangles represent regions of the MPC where the camera failed to take a photo. This stitched inspection montage provides a very useful way to view a large amount of inspection data at once and we can see features of interest such as scratches, corrosion, seams, etc. on the surface which span multiple inspection images are

[3]Different inspections are represented by the year and month, e.g. June 2020

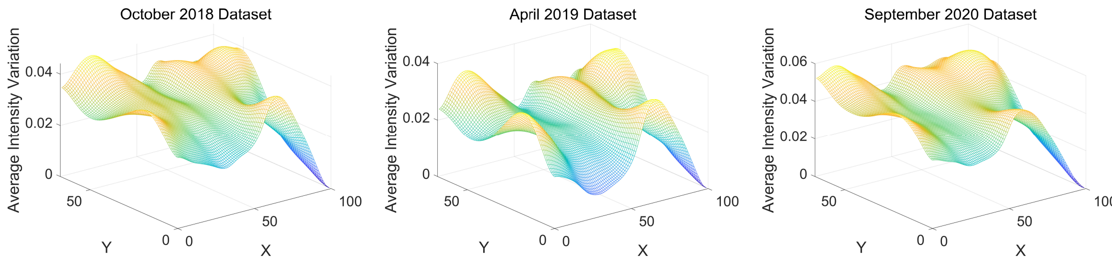

(a) Subset 1

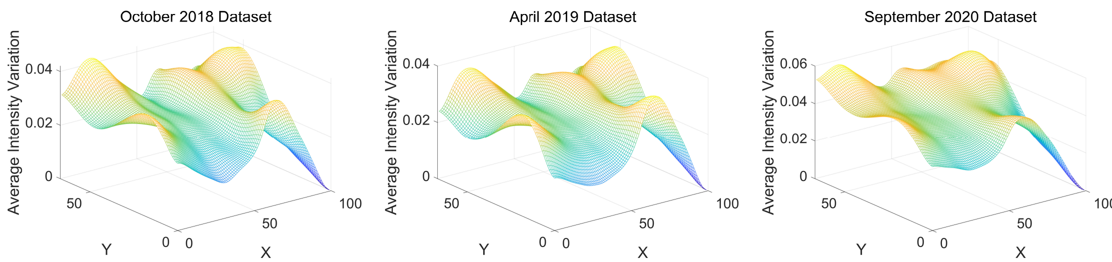

(b) Subset 2

Fig. 3: Two sets of illumination profiles calculated for each year. Each profile was calculated from 1000 images selected at random from a particular inspection dataset.

clearly represented.

Illumination Correction

To demonstrate the illumination correction algorithms we consider two raw inspection images. Given in Fig. 4 are two raw inspection images, Fig. 4a and 4c and the illumination corrected version of these images Fig. 4b and 4d. These light corrected images exhibit a marginal perceived improvement in quality but this is an important first step before attempting to correct other artifacts. Some inspection images which contain significantly more polished regions exhibit significant variation in appearance which is dependant on the orientation of illumination. Two examples are given in Fig. 5. The overlapping regions of these two images show the same region of text engraved on the surface of the MPC but appear significantly different. It can be clearly seen that physical features that are highlights in one image are shadows in a different image with different camera perspective. The shadow regions also exhibit a colour shift which makes correcting these effects even more challenging. Future work will look to correct these variations by first standardising the illumination for every inspection image.

(a) Original image

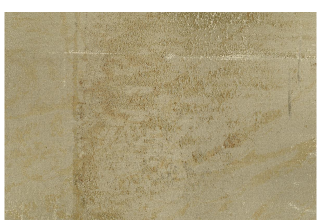

(b) Illumination corrected

(c) Original image

(d) Illumination corrected

Fig. 4: Two examples of illumination corrected images.

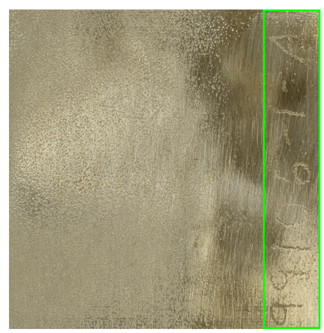

(a)

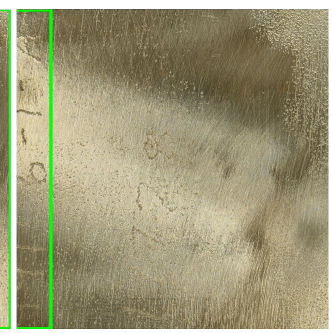

(b)

Fig. 5: Two inspection images which share a common region (highlighted in green) of the surface of the MPC.

CONCLUSIONS AND FUTURE WORK

In this paper we consider two problems related to the inspection of MPCs. The first is the problem of automatically stitching a very large dataset of images to generate a single

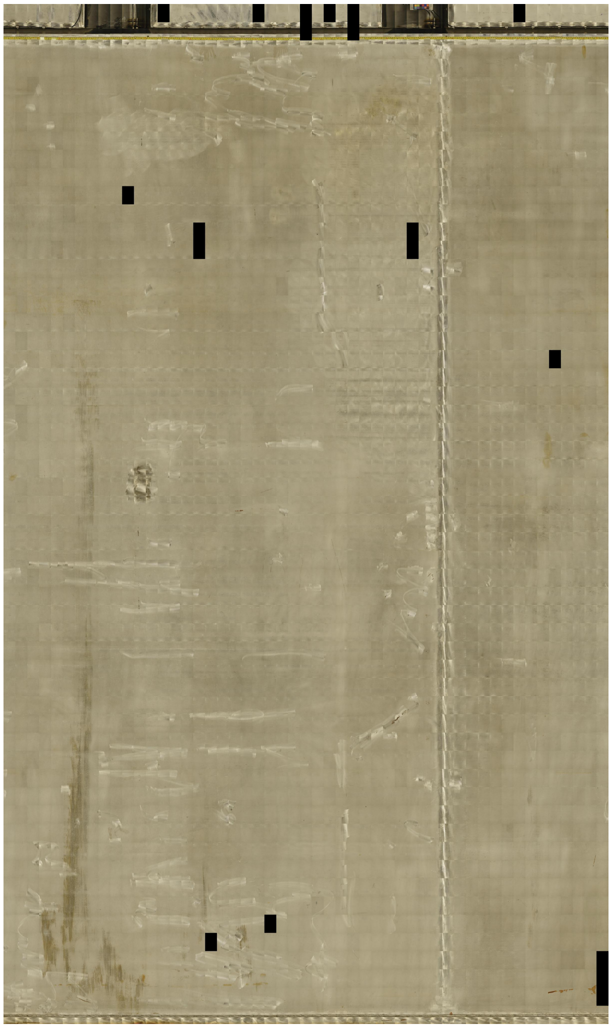

Fig. 6: Completed inspection montage after stitching has been performed.

inspection montage. We introduce a new set of algorithms for stitching MPC inspection footage. These algorithms performed very well and the initial results are very promising. The main features of interest on the MPC surface such as metal seams, corrosion, text on the metal, etc. are clearly resolvable over image boundaries. We then consider the problem of non-uniform illumination in the inspection images used to generate the montages. We present a novel method for the calculation of an illumination profile for the DSLR used in the camera rig and show that this is consistent between sub-sampled sets of our main datasets and also between multiple inspections recorded in different years. We outline a method for compensating for this variation in lighting and demonstrate it's application to a number of inspection images.

Finally we introduce some images where illumination correction alone isn't enough to improve appearance and there are still significant variations in the appearance of features such as text and this is very dependant on their position in any inspection image. This is due to the unique way that light interacts with the polished metallic surface and creates areas of highlight, areas of reduced lighting and a noticeable colour shift. Future work will focus on correcting these variations to allow the generation of montages where the inspection images are consistent in appearance and the appearance of features is independent of the camera orientation and position while preserving all the important fine detail.

REFERENCES

1. M. G. DEVEREUX, P. MURRAY, and G. WEST, "Automated object detection for visual inspection of nuclear reactor cores," *Nuclear Technology*, pp. 1–14 (2021).
2. 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**, 57–66 (2016).
3. Z. FEI, G. WEST, P. MURRAY, and G. DOBIE, "Automated generation of training dataset for crack detection in nuclear power plant components," in "12th Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (NPIC&HMIT 2021)," (2021).
4. S. GLASS, M. LEVESQUE, G. ENGELS, F. KLAHN, and D. FAIRBROTHER, "Underwater robotic tools for nuclear vessel and pipe examination," in "Proceedings ANS 8th International Topical Meeting on Robotics and Remote Systems," (1999), pp. 25–29.
5. M. DEVEREUX, P. MURRAY, G. WEST, S. BUCKLEY-MELLOR, G. COCKS, C. LYNCH, and A. FLETCHER, "Automated analysis of AGR fuel channel inspection videos," in "6th EDF Energy Nuclear Graphite Conference," (2018).
6. T. KHAN, M. KHAN, Y. KONG, and O. KITTANEH, "Stopping criterion for linear anisotropic image diffusion: a fingerprint image enhancement case," *EURASIP Journal on Image and Video Processing*, **2016** (02 2016).
7. T. TASDIZEN, E. JURRUS, and R. T. WHITAKER, "Non-uniform illumination correction in transmission electron microscopy," in "MICCAI Workshop on Microscopic Image Analysis with Applications in Biology," (2008), pp. 5–6.
8. W. A. MUSTAFA, H. YAZID, and M. JAAFAR, "An improved sauvola approach on document images binarization," *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, **10**, *2*, 43–50 (2018).
9. M. ESTER, H.-P. KRIEGEL, J. SANDER, X. XU, ET AL., "A density-based algorithm for discovering clusters in large spatial databases with noise." in "KDD," (1996), vol. 96, pp. 226–231.
10. D. GARCIA, "Robust smoothing of gridded data in one and higher dimensions with missing values," *Computational statistics & data analysis*, **54**, *4*, 1167–1178 (2010).
11. D. GARCIA, "A fast all-in-one method for automated post-processing of PIV data," *Experiments in fluids*, **50**, *5*, 1247–1259 (2011).