3-DvisualizationofAGRfuelchannelbricksusingStructure-from-Motion

KristoferLaw[a,1,*],GraemeWest[a],PaulMurray[a],ChrisLynch[b]

[a]*DepartmentofElectronicandElectricalEngineering,UniversityofStrathclyde,99GeorgeStreet,GlasgowG11RD,UnitedKingdom*[b]*EDFEnergyGeneration,BarnettWay,GloucesterGL43RS,UnitedKingdom*

Abstract

Thispaperoutlinesanewframeworkforapplying Structure-from-Motion(SfM)tochallenging,feature-poorenvironmentssuchasthoseobservedduringAGRfuelchannelinspection.DerivingstructuralinformationfromAdvancedGas-cooledReactor(AGR)inspectionfootageischallengingduetoseveralkeyissues:lackofdiscriminativesalientfeatureswithinthechannel,inconsistencyinlightingduringtheinspectionprocess,lackoftexturalinformationwithinthechannelandnoisefromtheinspectionequipment.Thispresentsdifficultiestotechniquessuchas SfMduetoitsrelianceonfindingandreliablytrackingasetofrobustfeaturesfrommultipleviewpoints.Thispaperintroducesthefirstuseofanincremental3-Dreconstructionframeworkwhichcanproducereconstructionsoffootageobtainedwithinanuclearreactor.Itapproachesthisissuebyintroducinganovelcorrespondencesearchingmethodologywhichcanoperatewithinfeature-poorenvironmentsbyutilisingaconstrained,iterativethresholdmatchingtechniquetoobtainrobustfeaturematches.Thispaperdemonstratestheapproachusingtwodatasets:laboratoryfootageobtainedfromanexperimentalsetupemulatingasmallsub-sectionofthechannelandin-coreinspectionfootageofAGRfuelchannels.

*Keywords:*3-DVisualization,Structure-from-Motion,AdvancedGas-cooledReactor

1.Introduction

1.1.Motivation

IntheUnitedKingdom,therearecurrently7AdvancedGas-cooledReactors(AGR)stationsinoperationwhichhaveexceededorareapproachingtheirinitial,conservative,designlifetimeestimations.Asthereactorages,itisimportanttomonitortheoperationaldegradationoftheAGRgraphitecore,anessentialcomponenttosafereactoroperationwhichisinherentlylife-limited.EachAGRgraphitecorecomposesofapproximately6000graphitebricks,around3000ofwhichformalatticestructureoffuelchannelswhichprovidemoderationandthestructureforhousinguraniumfuel,permittingCO_2coolantgasflowandhousingtheboroncontrolrodstoregulatethereaction.Duringreactoroperation,processessuchasfastneutronirradiationandradiolyticoxidationeffectivelyalter

*Correspondingauthor*Emailaddress:*kristofer.law@strath.ac.uk(KristoferLaw)

PreprintsubmittedtoNuclearEngineeringandDesign

thephysicalcharacteristicsofthegraphitebrickssuchasthestrength,weight,physicaldimensionsandYoung'smoduluswhichcaninducestructuraldefectssuchascracks[1].Consequently,itisimportanttomodel,measureandmonitorthereactorcoregraphitepropertiesandperformance,preeminentlyasthereactorcontinuestooperatebeyonditsinitialconservativedesignlifetime.

1.2.AGRInspection

ConditionMonitoring(CM)andRemoteVisualInspection(RVI)techniquesareemployedinparalleltoascertainthestructuralhealthoftheAGRreactorandtoensuresafeoperation.BothproceduresdifferwithCMapproachesbeingemployedonaregularbasiswithoutthenecessityoftemporaryreactorshut-down,quantifyingreactorhealthfromcontrolrodmotion[2]andFuelGrabLoadTrace(FGLT)data[3].RVIhoweverisdeployedduringscheduled,periodicreactoroutageswhereaportfolioofchannelsareselectedusingdecisionsupportsoftware[4]basedonuser-specificcriterion

(e.g. number of observed cracks or previous Television (TV) inspections) or the derivation of the optimal selection of fuel channels based on genetic algorithms using historical RVI and CM data. Visual Inspection is carried out by lowering tools such as the New In-Core Inspection Equipment (NICIE) or Channel Bore Inspection Unit (CBIU) into the fuel channels which obtain video footage. Video footage is gathered in three distinct stages: down channel footage where the camera is lowered to the bottom of the fuel channel facing downwards, allowing for engineers to view the fuel channel on the periphery and inspect the debris pot at the bottom of the channel. The second stage is the channel wall footage where a mirror is engaged to change the observable orientation of the camera onto the channel wall. The suspended camera vertically scans the fuel channel interior at 6 orientations $\theta = \{0°, 60°, 120°, 180°, 240°, 300°\}$ with an overlap of $\pm 5°$ between scans to provide full, circumferential coverage of the fuel channel interior. The final stage of inspection is where engineers will revisit sections of the channels which contain irregularities or known defects to capture crack following footage to allow for a comprehensive analysis of the perceived region of interest.

To gain a greater understanding of the fuel channel condition, the video footage is rigorously analysed and the characteristics of the defects will be quantified. Currently, visualisation of the RVI footage is performed by producing manually stitched image montages of regions of interest such as cracks in the graphite brick. This technique however provides limited structural information about the entire fuel channel and the demand of improved comprehensive visualisation techniques that can function using pre-existing RVI footage and inspection methodologies has increased as the reactors continue to mature. Techniques such as [1] have improved the analysis of AGR fuel channels by producing panoramic images of the fuel channels termed "Chanoramas" however pre-existing analysis techniques provide limited structural information due to a lack of depth as a result of the image capture process. To ascertain structural information, image processing techniques such as Structure-from-Motion (SfM), which can derive 3-D information of a scene given 2-D images from multiple views has been introduced to the AGR fuel channels in previous papers by the authors [5, 6] at a conceptual level - highlighting the difficulties with applying generic SfM frameworks in challenging environments. This paper proposes an application-specific SfM framework which utilises pre-existing inspection data to reproduce representative sparse 3-D reconstructions of vertical scans of AGR graphite brick interior surface.

1.3. Contribution

The contributions made in this paper are:

1. Development of a novel, application-specific SfM framework which can produce sparse 3-D reconstructions from RVI in-core inspection footage.
2. Introduction of a correspondence searching algorithm designed to disambiguate feature matching in feature-poor and challenging environments.
3. Evaluation of the framework using footage captured from novel inspection apparatus designed to emulate the inspection process of an AGR fuel channel.

In this paper, Section 2 outlines the current state of the art in nuclear reactor component visualisation. Section 3 introduces the SfM framework developed for 3-D visualisation of AGR fuel channels and Section 4 disseminates the key results of the 3-D AGR visualisation process before Section 5 bring the paper to a conclusion.

2. Related Work

2.1. Visualization within Nuclear

RVI and visualisation is deployed throughout the nuclear industry for a variety of tasks such as defect detection [1, 7, 8] and repair [9, 10], accident response [11] and reactor decommissioning [12] across many reactor designs such as the AGR, the Canada Deuterium Uranium (CANDU) reactor, Boiling Water Reactor (BWR) and the Pressurized Water Reactor (PWR). With each reactor having unique architectural characteristics, application-specific equipment is often developed to visually inspect a variation of life-limited components pertinent to each aforementioned reactor design.

A variation of RVI methodologies are deployed such as Ultrasonic tools to investigate weld joints [13] in the Reactor Pressure Vessel (RPV) and defects in CANDU pressure tubes [7], bespoke underwater robotic camera systems to investigate the

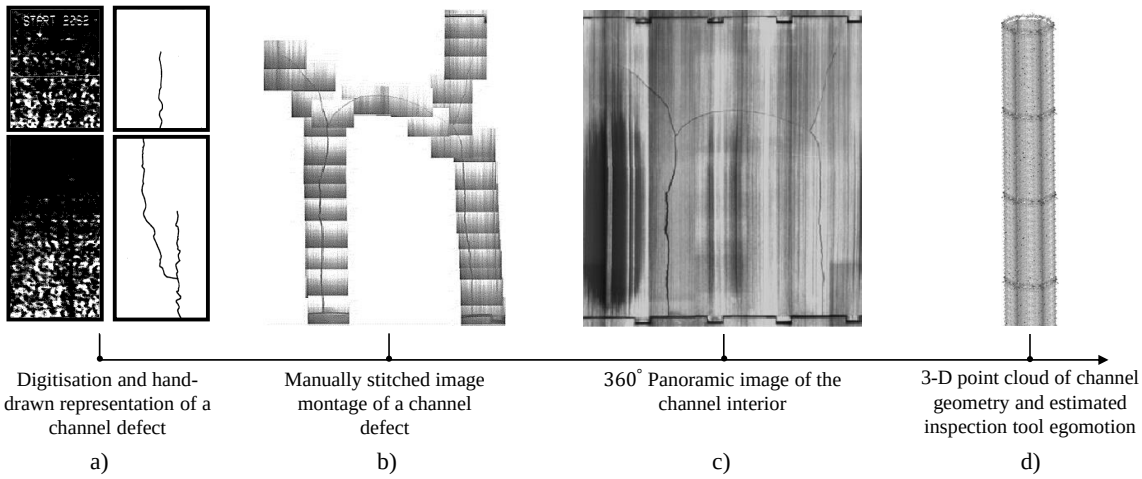

Figure 1: Timeline of how visualisation techniques have evolved over time from manual hand-drawn interpretations to 3-D reconstructions of RVI footage.

PWR RPV [10] and lower core plates [9], pipe inspection robots for feeder pipes in the Pressurized Heavy Water Reactor (PHWR) and tethered inspection tools used within the AGR fuel channels and control rods [2, 3, 1]. The majority of visualisation techniques however are limited to using signal and 2-D image processing to evaluate components within the respective reactor designs, with 3-D visualization techniques only being recently deployed to evaluate Fukushima [11, 12] to assess reactor station damage.

2.2. AGR Visualization - A Timeline

AGR visualization was initiated approximately 25-30 years ago not long after AGR start of life where binary digital image stills were extracted from analog tape and hand-drawn representations were produced of defects [6]. This process however was time-consuming due to the digitisation of the analog video tape and provided very limited means of visualization in regards to the condition of the AGR fuel channel. As the reactors age, inspection tools such as the CBIU and NICIE1/2 were introduced to deliver improved image quality and subsequently provided the basis for the modern visualization methodologies currently utilised to inspect AGR fuel channels which are manually stitched montages of Regions of Interest (ROI) and the "chanoramas" which are 360° unwrapped panoramic images of the AGR fuel channels produced by the Automated Software Image Stitching Tool (ASIST) [1]. The final perceived stage of AGR visualization is the use of so called pivot videos (i.e. a video which produces an illusionary stereoscopic effect by pivoting around an ROI) [14] and the preliminary application of SfM [5, 6] to derive structural depth information with pre-existing RVI footage. An approximate time-line for the evolution of AGR visualization can be seen in Fig.1.

2.3. Application specific difficulties of 3-D Reconstruction

There are several key challenges when applying SfM to AGR RVI footage which can be broadly attributed to two sources: 1) environmental influence of how the inspection tool operates and 2) tool-specific issues which occur due to the inspection tool not being inherently designed for 3-D visualization [6]. First of all, the environment in which the inspection tool operates is highly constrained and the graphite brick surface often lacks discriminative salient features in addition to repetitive textures within the channel which causes issues with reliably obtaining robust correspondences between images. With regards to the inspection tool, there is inconsistency in lighting during the inspection process due to engineer intervention, noise from the camera and furthermore due to the camera being tethered, it is subject to undesirable rotation and abrupt translations when interfacing with the brick surfaces. These issues combined together makes it extremely challenging for generic SfM frameworks to produce representative 3-D structural geometry of the AGR fuel channels so a bespoke framework must be designed to alleviate these issues.

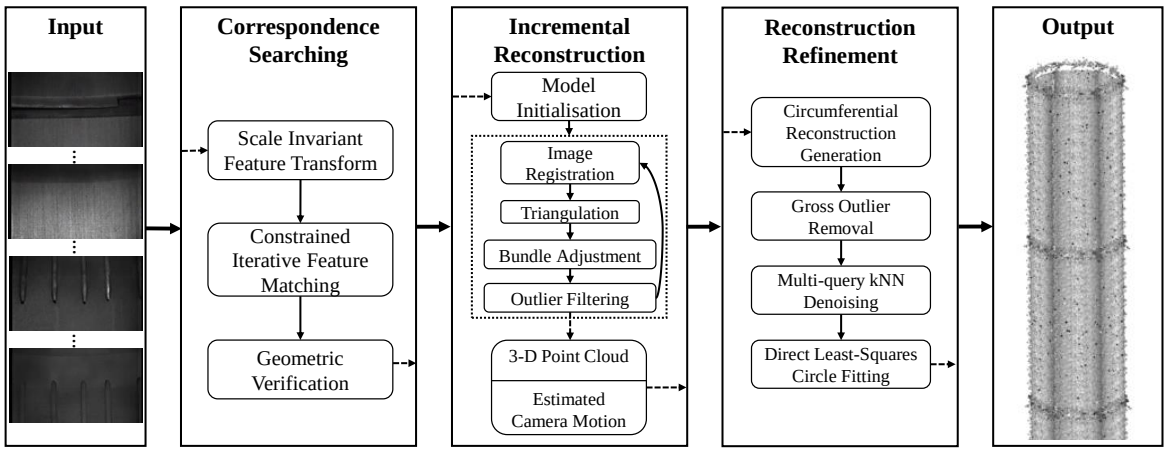

Figure 2: Overview of the 3-D reconstruction framework developed for the visualization of AGR graphite brick interiors.

3. AGR Reconstruction Framework

The 3-D reconstruction process is decomposed into three key stages: the first stage comprises of pre-processing the RVI inspection videos (removal of inspection tool overlay) and decomposing them into images. The second stage performs correspondence searching between the input inspection images for each orientation and produces a robust candidate set of feature correspondences. The tertiary stage takes the ascertained correspondences and incrementally reconstructs a sparse 3-D point cloud. The final stage takes the reconstructions for each orientation, combining them into a singular, cohesive cylindrical structure before de-noising and refining the final 3-D point cloud. This process is illustrated in Fig 2.

3.1. Correspondence Searching

Generic correspondence searching can be segregated into three sub-tasks: feature detection, feature matching and geometric verification. Feature detection is the first process where images within the image dataset are reduced to an abstract, sparse representation based on salient features in the image such as edges and blobs in the form of appearance descriptors. Each descriptor within the image describes the surrounding neighbourhood characteristics of detected, salient key-points within the image which gives each key-point scale and illumination invariant properties. Feature matching occurs afterwards where the descriptors produced for each image is matched with subsequent images and their corresponding descriptors using a similarity or distance-based measure to ascertain correspondences. The geometric verification stage then evaluates the detected matches under the assumption that all feature points go under a similar geometric transformation and discards matches that do not conform. Due to the issues previously described in Section 2.3, the generic SfM framework however does not provide adequate results therefore the correspondence searching stage must be approached in a different manner.

The correspondence searching mechanism deployed within our framework is more rigorous in finding stronger and more robust matches at the expense of increased computation. The framework deployed first utilises the Scale Invariant Feature Transform (SIFT) [15] to generate uniformly sampled appearance descriptors for keypoints within each image in the dataset. Subsequently, the image feature sets $\mathcal{F}_i = \{(\mathbf{x}_j, \mathbf{f}_j) | j = 1, ..., N_{F_i}\}$ where $\mathbf{f}_j$ is the SIFT appearance descriptor and $\mathbf{x}_j = [x, y] \in \mathbb{R}^2$ is the detected feature locations are bi-directionally matched in a sequential manner, deriving the squared euclidean distance $d = \sqrt{\sum_{i=1}^{N_{F_i}} (\mathcal{F}_a - \mathcal{F}_b)^2}$ between the respective feature sets in I_a and I_b to obtain a set of unique matches $\mathcal{M}_{ab} \in \mathcal{F}_a \times \mathcal{F}_b$ where $\mathcal{M}_{ab} = \text{argmin}_{\mathbf{f}_j \in \mathcal{F}} d(\mathcal{F}_a, \mathcal{F}_b)$.

After a candidate set of putative matches have been obtained, in feature-poor environments such as the mid-brick layer, the detected number of matches $N_{matches}$ is often insufficient for a comprehensive representation of the 3-D structural geometry or in extreme cases, insufficient matches for the estimation of rigid body motion. Relaxation tech-

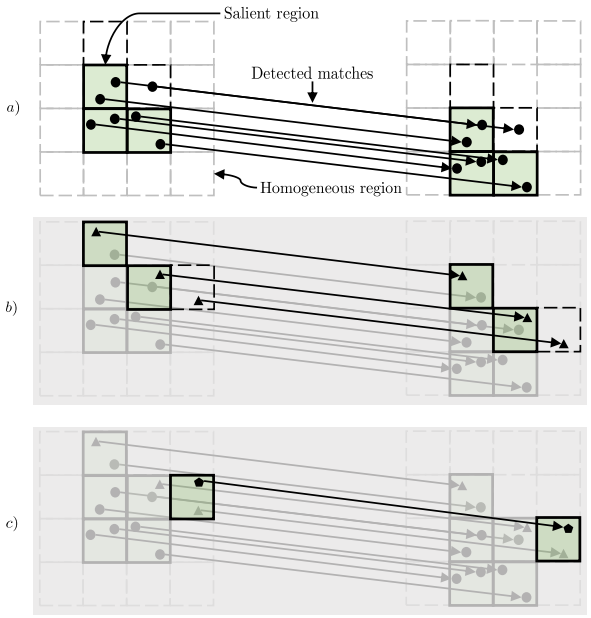

Figure 3: Visualisation of the feature matching methodology between two contiguously captured coherent images, matching at varying thresholds and identifying strong/poor neighbourhoods and feature points. a) The principal iteration identifying the saliency of the block regions; b)&c) Demonstrating further iterations and identification of new feature matches through iterative threshold relaxation.

niques or the use of re-triangulation is often used [16, 17] to incorporate a larger amount of matches. However, this results in the potential integration of ambiguous matches which can result in inaccurate feature tracking and increased error drift in incremental reconstructions. To avoid these issues as identified by [18], we take into account the spatial distribution of the detected feature matches using an iterative approach as visualised in Fig 3. An initial matching procedure takes place based on the Best Bin First (BBF) matching principle [19] where the ratio of distances d between the closest match and second closest match multiplied by t is not greater than the ratio between all other feature descriptors for the image. With distance threshold t being set high, ensuring very strong matches in the primary iteration. Once the principal iteration has been completed, the movement can be initially verified using two distinct approaches: the use of a-priori knowledge or the use of dense motion estimation techniques.

Implementing translation-based geometric constraints established using a-priori knowledge about the rigid body motion is the simplest method of disambiguating candidate matches between image pairs. The NICIE/CBIU inspection camera moves approximately vertically ± 5 pixels between each image decomposed at the camera frame rate of 25 frames per second (FPS) [1]. Therefore, a simple translation filter is implemented to ensure the feature locations $\mathbf{x}_j$ for each identified match fall within a boundary constrained by the horizontal (α) and vertical motion (β) so that for each corresponding point, $\exists!\ \mathbf{x}_j = [x \pm \alpha, y \pm \beta]$. Another method of capturing the motion between frames is the use of optical flow, a technique which characterises the dynamic translational motion between contiguously captured images and returns an sparse [20, 21] or dense estimation of the motion [22] in the form of translation vectors $[x, y]$. Within our framework, we extract the underlying image regions pertaining to the identified *strong* feature match regions in I_n and utilise the dense motion algorithm Dual-Total Variation L^1 norm (Dual TV-L^1) [22] for its ability to handle changes in illumination and suppress noise to estimate the dense regularized motion $\mathcal{O}$ within the regions. To phase out potential outliers, we obtain the median translation $[\tilde{x}, \tilde{y}]$ and this approach offers a higher degree of flexibility with the tool rigid body motion often being inconsistent, subject to unwanted rotation or engineer intervention during inspection [6].

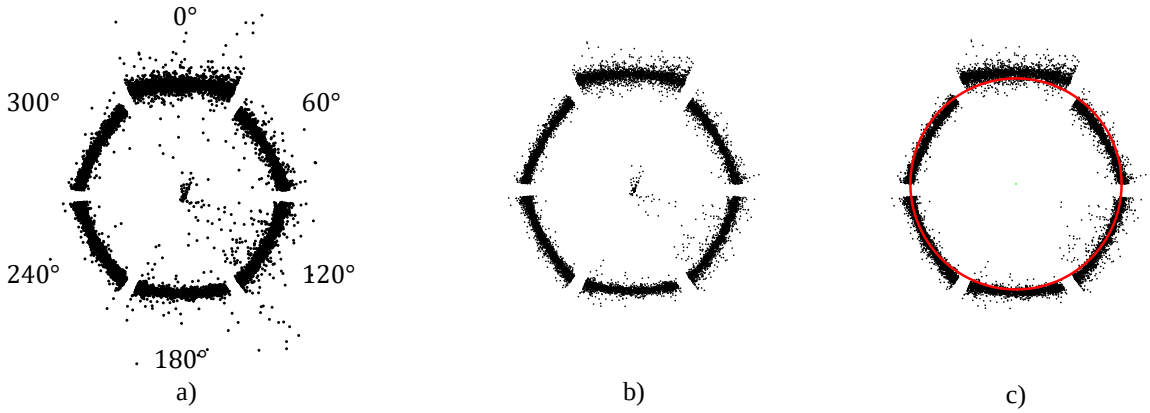

Figure 4: Top-down view of the 3-D points at each stage of refinement. a) contains the combined, original 3-D point cloud; b) A point cloud denoised by multi-query kNN denoising; c) Further point-cloud denoising using circle fitting.

If $|\mathcal{M}_{ab}| \geq N_{matches}$ and the matches are verified, no more iterations are carried out. However, if $|\mathcal{M}_{ab}| \leq N_{matches}$ on the principal iteration, the framework will continually relax t until $N_{matches}$ is exceeded and the verification process is integrated to remove ambiguous mismatched candidates. If the iterative process fails to provide an adequate number of matches, bounded Zero Mean Normalised Cross-Correlation (ZNCC) is applied or the image pair is discarded from the reconstruction.

3.2. Incremental Reconstruction

During RVI inspection of the AGR fuel channel, each camera view has a very limited perspective of the entire AGR fuel channel at any given time which means that the resulting 3-D representation of the channel must be cumulatively generated as the camera moves either up or down within the AGR fuel channel. The 3-D reconstruction process can be implemented in three modes: Global, Hierarchical and Incremental which considers all views, subsets of views or evaluates the views in a serial fashion. Due to the repetitive, pseudo-cyclic environment of the AGR fuel channel interior, global and hierarchical methods are inadequate and therefore an incremental approach is admissible. As demonstrated in the first stage of Fig 2, the model must be first initialized using a two-view reconstruction and for incremental methods, strong initialisation is critical [17] to minimise error drift and improve the reconstruction quality. To ensure the best possible initialization, $N_{matches}$ required for the first initial two-view reconstruction is increased to 350 matches. After initialisation, each image is registered and the feature matches determined by the correspondence searching stage is triangulated in a 3-D space. Once the triangulation has taken place, an error minimization process is undertaken to reduce the reprojection error between the 2-D feature location $\mathbf{x}_j$ and the projected 3-D point location $\mathbf{X}_j$ which in turn refines $\mathbf{X}_j$ and the corresponding camera pose. This process is repeated for each orientation during the inspection process.

3.3. Reconstruction Refinement

The final module of the reconstruction framework is used to bring together all the orientation-based reconstructions into a refined, cohesive circumferential 3-D reconstruction that is representative of the AGR fuel channel. As seen in Fig 2, the first stage is to rotate the 3-D points of each orientation point-cloud relative to the captured RVI footage orientations θ. Each orientation point cloud p_θ and the associated 3-D points $\mathbf{X} \equiv (\acute{x}, y, \acute{z})^T$ are rotated around the y-axis where $\acute{x} = x\cos\theta + z\sin\theta$ and $\acute{z} = z\cos\theta - x\sin\theta$ and merged into a singular 3-D point cloud.

Afterwards, the refinement process is performed using an approach comprising of three sequential stages: gross-outlier removal, multi-query k-Nearest Neighbours (kNN) point cloud denoising and the use of a direct least-squares circle fitting technique to remove noisy outliers. Gross outlier removal is the first stage, which combines the use of a filter that removes 3-D points which fall outside a hard-coded boundary and the Median Absolute Deviation (MAD) to remove gross-outliers. The second stage applies the multi-query kNN technique

devised by [23] which queries the mean distance of all 3-D points within the model to each neighbouring point and which allows for the detection and removal of outliers. The final stage integrates the start-of-life structural parameters of the AGR channel in the form of robust circle fitting onto a 2-D co-ordinate space comprising of $(\acute{x}, \acute{z})$ components of each 3-D point in order to further remove outliers from the model. Once the outliers have been removed, each partition pertaining to the points of each orientation are translated to produce a single, cohesive cylindrical 3-D reconstruction.

4. Results

4.1. Datasets & Evaluation Procedures

To test the proposed framework, real in-core RVI footage obtained from multiple reactors and laboratory footage captured from an experimental apparatus comprising of three interlocked graphite bricks, two of which are intentionally cracked are used to test and validate the framework performance respectively. The laboratory footage is acquired using a interior-facing camera with a spatial resolution of 1920×1080 at 30fps, downsampled to 720×405 to approximate the resolution of the AGR RVI footage. The camera was manually translated using a linear translation stage at a fixed distance inside the brick whilst simulating the inspection protocol observed within the RVI footage when performing individual orientation scans. To validate the orientational reconstruction performance of the proposed framework, a single scan of one layer of the experimental lab apparatus is used and the resulting reconstructions compared against two other excellent state-of-the-art methods [16, 17]. Circumferential reconstructions are performed using only the RVI in-core footage and using our own framework since this is more application specific and incorporates functionality not available in generalised SfM frameworks. To maintain parity with tested frameworks, feature detection/matching parameters are kept consistent and no geometric priors are used during correspondence searching[1].

4.2. Orientational Reconstructions

The laboratory footage used to validate the orientation-based reconstructions represents the best case scenario due to the image quality from the camera being sufficiently higher, lighting is globally consistent on an intra and inter-image basis and the camera movement is constrained to pure translation. The reconstruction results produced by our framework in addition to other incremental reconstruction frameworks are shown in Fig 5. As seen in Fig 5a) and d), the reconstruction results produced by the proposed framework correctly exhibits the observed structure with the point-cloud adhering to a solid arc-shape (this shape can also be seen in Fig 4 where each orientation reconstruction follows an arc) and in Fig 5d) where the camera frustum remains at a fixed distance whilst maintaining minimal deviation from the expected pure translation. Contrarily, the point cloud reconstructions in b/e) and c/f) belonging to COLMAP and VisualSFM clearly demonstrate high degrees of distortion with the point-cloud and the corresponding camera paths bending both forwards and backwards - indicating a large degree of cumulative drift predominant within incremental systems [18]. With regards to point cloud density (i.e. the completeness or resolution of the point-cloud), VisualSFM and COLMAP vastly outperforms our framework with 144,331 and 88,829 3-D points whereas the proposed framework produces around 5494 3-D points. This is a direct result of the proposed framework searching for stable and robust feature matches at the expense of reconstruction sparsity. The time complexity of the correspondence searching procedure is $O(n^2)$ and the incremental reconstruction process $O(n^4)$ where n is the number of images. To speed up this procedure, the image data is manually segmented based on scan and orientation, and reconstructions are performed of the requested orientation or brick layer to minimise the number of images. Due to the conceptual nature of the framework, unlike VisualSFM or COLMAP, it does not take advantage of GPU-accelerated feature matching or multicore bundle adjustment algorithms [16] and this is an area of future work.

4.3. Circumferential Reconstructions

Circumferential reconstructions are performed using RVI in-core inspection footage and an example of these reconstructions can be seen in Fig 6. With circumferential reconstructions, the structure of the brick(s) of interest can be representatively reconstructed in a sparse manner given that the footage of all orientations is of sufficient quality. With these reconstructions, the fiducial fea-

[1]It is also important to note that the camera frustum produced by our framework is down-sampled for better visibility.

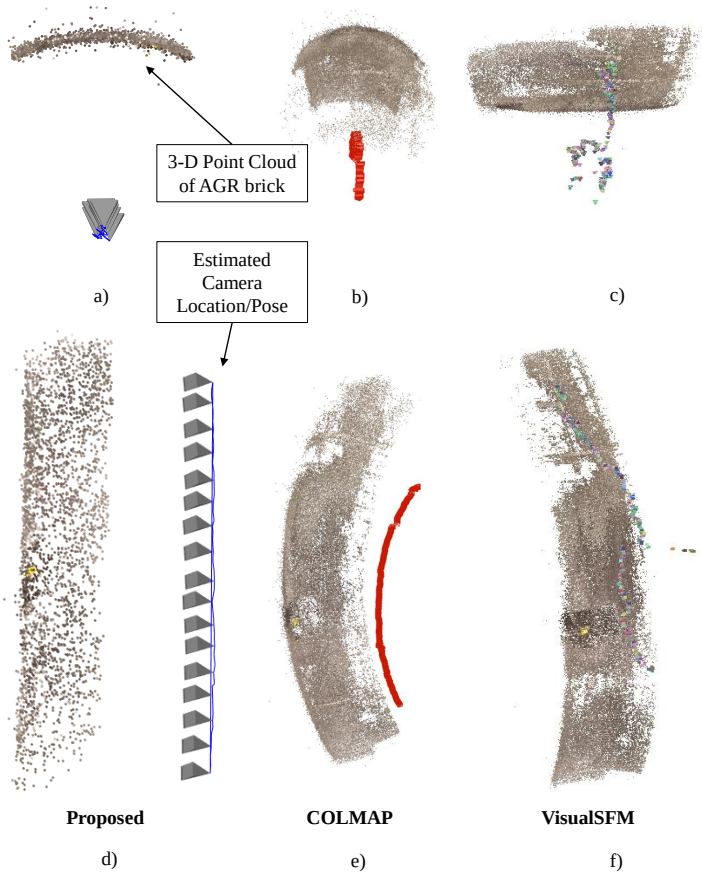

Figure 5: Results obtained using AGR laboratory footage. a-c) Above camera view of the AGR brick point cloud reconstructions with the associated camera trajectory using our reconstruction framework, COLMAP and VisualSFM respectively; d-f) Camera view of a point cloud reconstruction of a scan obtained within a single AGR brick using our reconstruction framework, COLMAP and VisualSFM respectively.

tures of the bricks can be identified such as the mid-brick layer and brick-interfaces between each brick. However, since the circumferential generation is inherently post-processing of the point-cloud, when merging the reconstructions, there is no automatic procedure currently in place to match and merge the 3-D points within the point cloud due to the difficulty of inter-scan matching as a result of the small 10° overlap.

5. Conclusion

In this paper, we have presented a method for extracting 3-D structural geometry in the form of a sparse 3-D point cloud from RVI footage. The proposed framework extends upon previous visualisation methods and provides structural information which can be useful for inspection engineers to gain a deeper understanding of the structural health of the AGR fuel channels whilst exploiting pre-existing or historical AGR RVI footage. The framework utilises a strict, motion guided matching scheme to extract strong, robust correspondences between images within a challenging environment which allows for the reproduction of representative 3-D point clouds of the AGR fuel channels and evaluates the framework using both experimental and in-core RVI footage. The proposed framework has the capability to not only be applied to AGR fuel channels but also other footage which observes demanding, low feature environments. The use of this framework in parallel with the ASIST software de-

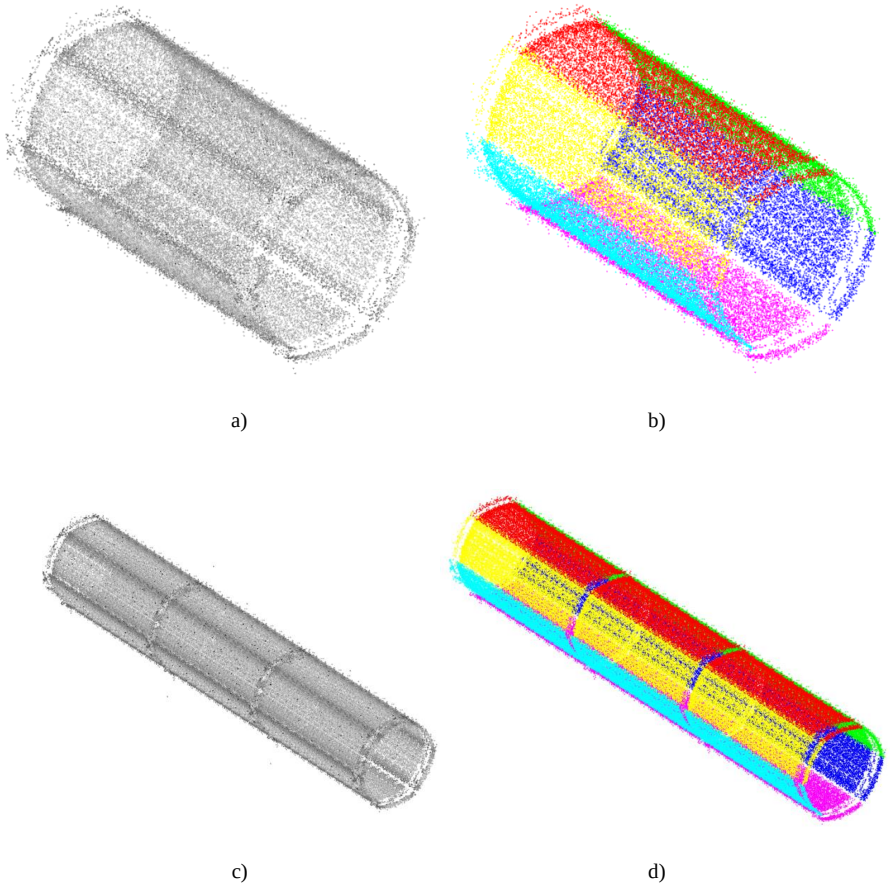

Figure 6: Results obtained using in-core footage. a) Circumferential brick layer; b) Colored (Each orientation having a unique color) circumferential brick layer; c) Reconstruction consisting of three layers; d) Colored three layer reconstruction

veloped by [1] and the "pivot" videos [3] visualisation mechanisms aims to extract sizeable amounts of actionable information for diagnostic and prognostic purposes whilst utilising pre-existing AGR RVI footage.

Acknowledgements

The authors would like to thank EDF Energy for funding this work.

References

[1] P. Murray, G. West, S. Marshall, S. McArthur, Automated in-core image generation from video to aid visual inspection of nuclear power plant cores, Nuclear Engineering and Design 300 (2016) 57–66. doi:10.1016/j.nucengdes.2015.11.037.
URL http://dx.doi.org/10.1016/j.nucengdes.2015.11.037

[2] C. J. Wallace, G. M. West, G. J. Jahn, S. D. J. McArthur, D. Towle, G. Buckley, Control Rod Monitoring Of Advanced Gas Cooled Reactors, in: Seventh American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies, ANS, Las Vegas, Nevada, USA, 2010, pp. 1–10.

[3] G. M. West, C. J. Wallace, S. D. McArthur, Combining models of behaviour with operational data to provide enhanced condition monitoring of AGR cores, Nuclear Engineering and Design 272 (2014) 11–18. doi:10.1016/j.nucengdes.2013.12.067.
URL http://dx.doi.org/10.1016/j.nucengdes.2013.12.067

[4] C. E. Watson, P. C. Robinson, J. F. Burrow, P. R. Maul, The application of portfolio selection to fuel channel inspection in advanced gas-cooled reactors, Nuclear Engineering and Design 328 (January) (2018) 145–153. doi:10.1016/j.nucengdes.2018.01.013.
URL https://doi.org/10.1016/j.nucengdes.2018.01.013

[5] K. Law, G. West, P. Murray, C. Lynch, Towards Extracting 3-D Structural Representations of AGR Core

Fuel Channels from 2-D In-Core Inspection Videos, in: International Symposium on Future I&C for Nuclear Power Plants (ISOFIC 2017), KNS, Gyeongju, Republic of Korea, 2017, pp. 1–10.

[6] K. Law, G. West, P. Murray, C. Lynch, 3-D Advanced Gas-cooled Nuclear Reactor Fuel Channel Reconstruction using Structure-from-Motion, in: International Topical Meeting on Nuclear Plant Instrumentation, Control and Human Machine Interface Technologies, ANS, San Francisco, CA, USA, 2017, pp. 67–76.

[7] T. Lardner, G. West, G. Dobie, A. Gachagan, Automated sizing and classification of defects in CANDU pressure tubes, Nuclear Engineering and Design 325 (January) (2017) 25–32. doi:10.1016/j.nucengdes.2017.09.029.
URL http://dx.doi.org/10.1016/j.nucengdes.2017.09.029

[8] F.-C. Chen, R. M. R. Jahanshahi, NB-CNN: Deep Learning-based Crack Detection Using Convolutional Neural Network and Naïve Bayes Data Fusion, IEEE Transactions on Industrial Electronics 65 (5) (2017) 4393–4400. doi:10.1109/TIE.2017.2764844.
URL http://ieeexplore.ieee.org/document/8074762/

[9] B. H. Cho, S. H. Byun, C. H. Shin, J. B. Yang, S. I. Song, J. M. Oh, KeproVt: Underwater robotic system for visual inspection of nuclear reactor internals, Nuclear Engineering and Design 231 (3) (2004) 327–335. doi:10.1016/j.nucengdes.2004.03.012.

[10] J. H. Kim, J. C. Lee, Y. R. Choi, LAROB: Laser-guided underwater mobile robot for reactor vessel inspection, IEEE/ASME Transactions on Mechatronics 19 (4) (2014) 1216–1225. doi:10.1109/TMECH.2013.2276889.

[11] K. Nagatani, S. Kiribayashi, Y. Okada, K. Otake, K. Yoshida, S. Tadokoro, T. Nishimura, T. Yoshida, E. Koyanagi, M. Fukushima, S. Kawatsuma, Emergency Response to the Nuclear Accident at the Fukushima Daiichi Nuclear Power Plants using Mobile Rescue Robots, Journal of Field Robotics 30 (1) (2013) 44–63. doi:10.1002/rob.21439.

[12] A. Seki, O. J. Woodford, S. Ito, B. Stenger, M. Hatakeyama, J. Shimamura, Reconstructing Fukushima: A Case Study, in: 2014 2nd International Conference on 3D Vision, IEEE, Tokyo, Japan, 2014, pp. 681–688. doi:10.1109/3DV.2014.75.
URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7035885

[13] P. Nanekar, N. Jothilakshmi, T. Jayakumar, Ultrasonic phased array examination of circumferential weld joint in reactor pressure vessel of BWR, Nuclear Engineering and Design 265 (2013) 366–374. doi:10.1016/j.nucengdes.2013.09.015.
URL http://dx.doi.org/10.1016/j.nucengdes.2013.09.015

[14] G. West, P. Murray, S. Marshall, S. Mcarthur, Improved Visual Inspection of Advanced Gas-Cooled Reactor Fuel Channels, International Journal of Prognostics and Health Management.

[15] D. G. Lowe, Distinctive image features from scale-invariant keypoints, International Journal of Computer Vision 60 (2) (2004) 91–110. doi:10.1023/B:VISI.0000029664.99615.94.

[16] C. Wu, Towards linear-time incremental structure from motion, in: International Conference on 3D Vision, 3DV 2013, IEEE, Seattle, Washington, USA, 2013, pp. 127–134. doi:10.1109/3DV.2013.25.

[17] J. L. Schönberger, J.-M. Frahm, Structure-from-Motion Revisited, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Las Vegas, Nevada, USA, 2016, pp. 4104–4113. doi:10.1109/CVPR.2016.445.

[18] S. El Kahi, D. Asmar, A. Fakih, J. Nieto, E. Nebot, A vison-based system for mapping the inside of a pipe, in: International Conference on Robotics and Biomimetics, no. 1, 2011, Phuket, Thailand, 2011, pp. 2605–2611. doi:10.1109/ROBIO.2011.6181697.

[19] J. S. Beis, D. G. Lowe, Shape indexing using approximate nearest-neighbour search in high-dimensional spaces, in: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97), CVPR '97, IEEE Computer Society, Washington, DC, USA, 1997, pp. 1000–.
URL http://dl.acm.org/citation.cfm?id=794189.794431

[20] S. Baker, I. Matthews, R. Patil, G. Cheung, I. Matthews, R. Gross, T. Ishikawa, I. Matthews, Lucas-Kanade 20 Years On: A Unifying Framework, International Journal of Computer Vision 56 (3) (2004) 221–255. doi:10.1023/B:VISI.0000011205.11775.fd.
URL http://www.springerlink.com/openurl.asp?id=doi:10.1023/B:VISI.0000011205.11775.fdhttp://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.72.9334{&}rep=rep1{&}type=pdf

[21] J. Maier, M. Humenberger, M. Murschitz, O. Zendel, M. Vincze, Guided Matching Based on Statistical Optical Flow for Fast and Robust Correspondence Analysis, in: M. Leibe, Bastian and Matas, Jiri and Sebe, Nicu and Welling (Ed.), European Conference on Computer Vision (ECCV), no. 1, Springer International Publishing, 2016, pp. 101–117. arXiv:1311.2901, doi:10.1007/978-3-319-46448-0.
URL http://www.nature.com/doifinder/10.1038/nature14539{%}5Cnhttp://www.nature.com/doifinder/10.1038/nmeth.3707{%}5Cnhttp://link.springer.com/10.1007/978-3-319-46448-0

[22] A. Wedel, T. Pock, C. Zach, H. Bischof, D. Cremers, An improved algorithm for TV-L 1 optical flow, in: Statistical and Geometrical Approaches to Visual Motion Analysis, Vol. 5604, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 23–45. doi:10.1007/978-3-642-03061-1_2.
URL http://link.springer.com/chapter/10.1007/978-3-642-03061-1{_}2

[23] R. B. Rusu, Z. C. Marton, N. Blodow, M. Dolha, M. Beetz, Towards 3D Point Cloud Based Object Maps for Household Environments, Robot. Auton. Syst. 56 (11) (2008) 927–941. doi:10.1016/j.robot.2008.08.005.
URL http://dx.doi.org/10.1016/j.robot.2008.08.005