

Reconstruction of 3D Concrete Microstructures Combining High-Resolution Characterization and Convolutional Neural Network for Image Segmentation



A. Cheniour
A. Ziabari
E. Tajuelo Rodriguez
M. Alnaggar
Y. Le Pape
T. M. Rosseel

June 2022



DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via OSTI.GOV.

Website: www.osti.gov/

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: 703-605-6000 (1-800-553-6847)
TDD: 703-487-4639
Fax: 703-605-6900
E-mail: info@ntis.gov
Website: <http://classic.ntis.gov/>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone: 865-576-8401
Fax: 865-576-5728
E-mail: report@osti.gov
Website: <https://www.osti.gov/>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Light Water Reactor Sustainability Program Milestone Report: M3LW-22OR0403013

**RECONSTRUCTION OF 3D CONCRETE MICROSTRUCTURES COMBINING
HIGH-RESOLUTION CHARACTERIZATION AND CONVOLUTIONAL NEURAL
NETWORK FOR IMAGE SEGMENTATION**

A. Cheniour
A. Ziabari
E. Tajuelo Rodriguez
M. Alnaggar
Y. Le Pape
T. M. Rosseel

June 2022

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	vii
ABBREVIATIONS	ix
ACKNOWLEDGMENTS	xi
EXECUTIVE SUMMARY	1
1 INTRODUCTION	2
2 3D CHARACTERIZATION OF CONCRETE SPECIMENS	4
3 SEMANTIC SEGMENTATION OF XCT-BASED IMAGES	6
3.1 Application of Traditional Semantic Segmentation Methods	6
3.2 Deep Learning-Based Semantic Segmentation	7
3.3 U-Net Architecture	7
3.4 Network Training	10
3.5 Results and Performance Metrics	11
3.6 Enhancing Accuracy Through Virtual Concrete Mesostructures	14
4 APPLICATION TO IRRADIATED CONCRETE MODELING	18
4.1 3D Simulation Setup	18
4.2 Results	20
5 DISCUSSION	22
6 CONCLUSION	23
7 REFERENCES	24

LIST OF FIGURES

1	Cored and sliced cylinders of Con-A and Con-B [1].	4
2	Calculated transmission for both concrete formulations in dry and saturated states for thicknesses of 10 to 40 mm.	5
3	Example original XCT slice of the concrete specimen.	8
4	Application of unsupervised clustering algorithms to segment the image shown in Figure 3.	8
5	Schematic of the proposed 2.5D U-Net architecture.	9
6	Manual labeling of slice 210 shown in Figure 3.	11
7	Segmented XCT images using U-Net.	12
8	Clipped stack of labeled images (slices 200 to 850) using U-Net. Background voxels (class 0) were removed in this illustration for visualization purposes.	13
9	Influence of the number of cuts (N) and the elongation on the aggregate shape. (reproduced from [2]).	15
10	Influence of resolution and number of cuts on the aggregate shape.	16
11	Influence of volume fraction (V_f) and resolution on the mesostructure.	16
12	Coring a cylinder ($h = 60\text{mm}$, $\phi = 40\text{ mm}$) from a 60 mm cube with $V_f = 0.605, 0.15$ mm voxels and 38–50 cuts per aggregate.	17
13	Reduced 3D simulation domain in MOSAIC (50 slices).	19
14	Phase map of the first layer of the 3D simulation domain.	19
15	Evolution of strain (%) in the xx, yy, and zz directions with fluence (n pm^{-2}) in quartz and hcp as predicted by MOSAIC.	20
16	Evolution of stress (MPa) in the xx, yy, and zz directions with fluence (n pm^{-2}) in quartz and hcp as predicted by MOSAIC.	21

LIST OF TABLES

1	Entered compositions for the transmission calculator for both concretes for 1 cm thickness and dry scenario, and resulting transmissions for thicknesses up to 40 mm for both saturated and dry conditions.	5
2	Precision, recall, and F1-score evaluation	14

ABBREVIATIONS

ANN	artificial neural network
ASR	alkali-silica reaction
CMG	Concrete Mesostructure Generator
CNN	convolutional neural network
DBSCAN	density-based spatial clustering of applications with noise
EMDA	expanded materials degradation analysis
FCN	fully convolutional network
FFT	fast-Fourier transform
GMM	Gaussian mixture model
hcp	hardened cement paste
HFIR	High Flux Isotope Reactor
IMAC	irradiated minerals, aggregate and concrete
IoU	intersection-over-union
IPI	interparticle interface
ITZ	interfacial transition zone
JCAMP	Japan Concrete Aging Management Program
LWRS	Light Water Reactor Sustainability
MOOSE	Multiphysics Object-Oriented Simulation Environment
MOSAIC	Microstructure-Oriented Scientific Analysis of Irradiated Concrete
MPI	message passing interface
NCT	neutron computed tomography
NPP	nuclear power plant
ORNL	Oak Ridge National Laboratory
RIVE	radiation-induced volumetric expansion
XCT	x-ray computed tomography

ACKNOWLEDGMENTS

This work was sponsored by the US Department of Energy Office of Nuclear Energy's Light Water Reactor Sustainability (LWRS) Program under contract DE-AC05-00OR22725 with UT Battelle LLC / Oak Ridge National Laboratory. The authors thank Shane C. Henderson (ORNL) and Robert A. Lefebvre (ORNL) for improving the parallel application of MOSAIC. The authors also thank Ercan Cakmak (ORNL) and J. David Arregui Mena (ORNL) for providing x-ray computed tomography (XCT) data of concrete specimens.

EXECUTIVE SUMMARY

After water, concrete is the second most used material in the world. Concrete's forming adaptability and low-cost constituents make it a predominant material used in the construction of civil infrastructures in nuclear power plants such as concrete biological shields, containment buildings, turbine buildings, fuel handling and storage buildings, underground piping for cooling, cooling towers, and so on. Depending on environmental and operating conditions, these passive structures are subject to time-dependent phenomena that can either enhance (e.g., continued hydration) or degrade concrete's structural performance. Unlike components such as the reactor pressure vessel and the primary circuit, concrete composition varies regionally because it is manufactured using local aggregates and cement. Hence, concrete performance metrics over time cannot be derived confidently using empirical relations. Alternatively, the specific characteristics of the local concrete constituents and their assemblage in concrete must be considered.

The US Department of Energy (DOE) Light Water Reactor Sustainability (LWRS) Program identified two major time-dependent mechanisms affecting the integrity of concrete and the associated structural performance in light water reactors: irradiation effects and alkali-silica reaction (ASR) [3]. Despite being either a physical (irradiation) or a chemical (ASR) degradation, both mechanisms originate in specific aggregate-forming minerals. Hence, the content and spatial distribution of these minerals have a considerable effect on the initiation and development of the induced damage. The dimensions of rock-forming minerals vary greatly between sub-micron and millimetric grains. In addition to the local effects caused by the composition and morphology of aggregates-forming minerals, the shape, dimensions, and distribution of the aggregates, which constitute about 70% of the concrete volume fraction, affect the formation, propagation, and percolation of cracks. Therefore, predictive models of aging concrete require adequate representation of the complex morphological features described above.

Successful efforts conducted by the LWRS Program have led to the development and validation of nonlinear mechanical models of irradiated concrete using the fast Fourier transform-based code MOSAIC informed by high-resolution multimodal 2D-imaging techniques. In light of the limitation of 2D representation to capture the residual mechanical properties of aged concrete, this report provides a high-accuracy methodology to reconstruct a 3D concrete microstructure, using a set of XCT-based images of a concrete specimen, to be used as a 3D simulation domain in MOSAIC. The approach described herein uses a modified version of the fully convolutional network (FCN) U-Net to label images and identify aggregates, cement paste, and pores or background based on a provided network training data set. The performance evaluation of the network revealed a successful application of this approach. An example 3D simulation of irradiated concrete in MOSAIC was performed using the generated 3D domain. This report also provides updates on 3D experimental characterization efforts. Finally, alternative approaches for image segmentation are discussed, as well as the computational limitations of 3D simulations in MOSAIC.

1 INTRODUCTION

Concrete is a complex, multiphasic, heterogeneous material made of natural aggregates (river gravel, crushed rock, and sand) and a binder formed by hydration of cementitious materials (Portland cement, fly ash, furnace slags, etc.). The cement paste contains pores with sizes ranging from a few nanometers up to approximately one millimeter. Moreover, after mixing and casting, air bubbles of a few millimeters may also be present. In a basic sense, concrete includes three constituent materials: it is a spread-out aggregate skeleton bound with cement paste containing porosity. Aggregates confer their stiffness to plain concrete: the ratio of Young's moduli of aggregate against the cement paste is approximately 4–5 in ordinary concrete. Hence, when subject to external loading, the mechanical stresses in concrete can be described as a flux following the aggregate's close packing and concentrated stress in the cement paste bridging aggregates. In general, the strength of aggregates is much higher than the strength of the cement paste, so early microcracking, fracture development, and propagation occur in the cement paste. Therefore, the morphological characteristics of the individual aggregates and the arrangements of the aggregates affect the fracture process and in turn affect the properties of the concrete. An accurate description and representation of the morphological characteristics of aggregates, air bubbles, and the cement paste are needed. However, the simplified 3-phase representation of concrete fails to describe the complexity of the damage caused by specific degradation mechanisms originating in the aggregates, such as the alkali-silica reaction and fast-neutron irradiation. These two mechanisms were identified as research priorities by the Expanded Materials Degradation Analysis (EMDA), a joint effort conducted by a panel of industry, regulators, academia, and national laboratories [3]. The EMDA guides the research activities of the US Department of Energy (DOE) Light Water Reactor Sustainability (LWRS) Program.

Alkali-silica reaction (ASR) is a term that was coined by Stanton in 1940 [4]. In the presence of highly alkaline solution (typically from high concentrations of sodium and potassium), poorly crystallized silica contained in the aggregates is gradually dissolved to form an alkali-calcic-silicate gel that expands from water absorption. The gel expansion results in pressure, thus causing microcracking. The location of gel formation and the associated induced cracking varies greatly from one type of aggregate to another: in embedded pockets, at the peripheries, along veins, and also in the cement paste porosity after transport. Therefore, the location of cracks primarily depends on the location of reactive aggregate-forming minerals such as chert, opal, and micro-quartz [5]. A precise description of the ASR's degradation ensures an accurate representation of the minerals' assemblage in the aggregates so that the reactive transport and poro-mechanics problems can be properly modeled.

Similarly, radiation effects on the concrete biological shield in nuclear power plants (NPPs) must be predicted to ensure safe long-term operations. Fast-neutron irradiation causes the gradual amorphization or *metamictization* of aggregate-forming minerals, leading to density changes, higher chemical reactivity, and changes to the mechanical properties of individual minerals. The extent of these changes varies according to the chemical properties of the minerals (i.e., composition, silica content, relative number of covalent to ionic bonds [6]). The differential radiation-induced volumetric expansion (RIVE) between minerals results in strain mismatch, thus causing the formation of inter- and intra-minerals cracks (see tentative models, [7, 8]). In addition, the overall RIVE of aggregates exerts significant deformation and pressure on the surrounding cement paste, resulting in dissipation mechanisms such as relaxation and cracking [9–11]. Therefore, an accurate model of the radiation-induced damage in concrete constituents and its effects on irradiated concrete properties requires a thorough description of the morphological characteristics of the minerals' assemblage in concrete aggregates.

Although different models have been developed at the microstructural scale [8, 12–14], they typically employ simplified representations of the concrete’s microstructure, such as a 2D domain, which assumes a uniformly elongated microstructure. Ideally, the state-of-the-art concrete radiation damage models would use a realistic 3D domain based on nondestructive characterization data to predict RIVE and damage development and propagation in the microstructure. The reconstruction of 3D concrete microstructures based on x-ray computed tomography (XCT) data for the purpose of modeling radiation damage is addressed in the previous milestone report [15]. In the latter, the reconstruction methodology employed an unsupervised learning method, namely the density-based spatial clustering of applications with noise (DBSCAN) [16], which is an advanced clustering algorithm implemented in the Scikit-learn Python library, to segment XCT-based images and then to eventually create a 3D domain for irradiated concrete simulations. In fact, conventional thresholding techniques failed to separate aggregates from the cement paste as a result of the low contrast between both phases. The combination of image processing and DBSCAN offered a significantly more accurate image segmentation output compared to other unsupervised learning techniques. Despite the overall good performance of the previous approach, several limitations are noted and discussed in the previous report, such as the need to reduce the image resolution for a reasonable computational time, unclear edges of identified clusters, and the limited identification of small aggregate clusters. It was concluded that convolutional network-based image segmentation would most likely lead to significant improvement in the results by using semantic segmentation architectures such as U-Net and SegNet.

This work applied a modified version of the convolutional network architecture U-Net [17] to perform image segmentation of XCT-based images of the Concrete-A sample provided by the Japan Concrete Aging Management Program (JCAMP) [1]. In this report, the current status of the 3D characterization efforts (XCT and neutron computed tomography (NCT)) are briefly summarized. Previously generated XCT data were used to train the network; this approach was applied on the full specimen’s volume. The performance metrics are presented herein, limitations of the proposed approach are discussed, and alternative approaches to generate 3D microstructures are recommended. The reconstructed concrete volume was employed as a 3D simulation domain in the fast Fourier transform-based (FFT) MOSAIC to simulate irradiated concrete. Finally, future needs for 3D MOSAIC simulations are discussed.

2 3D CHARACTERIZATION OF CONCRETE SPECIMENS

Cylindrical samples from Con-A and Con-B formulations [1] were cored using a core drill with an internal diameter of 2.8 cm and a height of 3.5 cm. The cored cylinders were cut in two halves along the vertical direction with a conventional diamond blade to ensure an appropriate sample thickness to yield above 30 per cent neutron transmission for NCT. The resulting samples are shown in Figure 1.



Figure 1. Cored and sliced cylinders of Con-A and Con-B [1].

XCT will be performed on the cut samples using a Zeiss Radia Versa 520 instrument with a resolution of $0.98 \mu\text{m}$. The samples will also be imaged with neutrons at the imaging instrument located in the High Flux Isotope Reactor (HFIR) at ORNL. The use of neutrons will be beneficial to enhance the contrast between aggregates and cement paste, because neutrons are sensitive to H atoms, and the cement paste is composed through multiple hydrated phases. Using the neutron transmission calculator provided by HFIR, estimations of neutron transmission in two different scenarios were performed for (a) full water saturation and (b) only chemically bound water in the cement paste (dry). The corresponding composition for the dry case is shown in Table 1. The resulting transmissions, shown in Figure 2, are similar for both cases. The samples were pre-dried to 62°C for several weeks, so their water content was estimated to lie between that of the dry and saturated cases. To obtain a successful tomography, transmission should be over 30 %. The maximum thickness that neutrons will need to penetrate in the cut samples is 1.4 cm, which results in transmission over 30% for both concrete formulations. A proposal for rapid access to the imaging beam line was submitted, and neutron beam time has been granted in early August.

Table 1. Entered compositions for the transmission calculator for both concretes for 1 cm thickness and dry scenario, and resulting transmissions for thicknesses up to 40 mm for both saturated and dry conditions.

Chemical formula	Thickness Con-A (mm)	Thickness Con-B (mm)	Density (g/cm ³)
H ₂ O	0.687813674	0.665262406	1
SiO ₂	5.936930405	5.257203915	2.65
Al ₂ O ₃	0.395177542	0.671489049	3.95
Fe ₂ O ₃	0.113762142	0.153044999	5.24
CaO	0.80291066	0.849018315	3.34
MgO	0.065504662	0.086210692	3.58
SO ₃	0.107640537	0.095763746	1.92
Na ₂ O	0.123575303	0.258033225	2.27
K ₂ O	0.086670893	0.118301917	2.35
TiO ₂	0.015530959	0.016336034	4.23
P ₂ O ₅	0.008235274	0.00939163	2.39
MnO	0.003277576	0.00196609	5.37

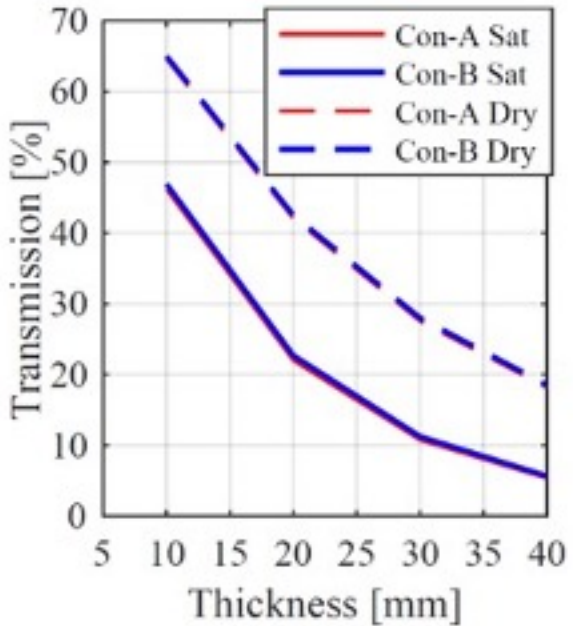


Figure 2. Calculated transmission for both concrete formulations in dry and saturated states for thicknesses of 10 to 40 mm.

3 SEMANTIC SEGMENTATION OF XCT-BASED IMAGES

Semantic segmentation is the process of labeling each data point in a data set (e.g., images) and attributing each data point to a specific class among a predefined set of classes. The nature of a class depends on the data type; in this case, the data set corresponds to the XCT concrete image, such as Figure 3, which shows the specimen's surrounding area (background), the aggregates surrounded by the cement paste, and the pores within the concrete. Ideally, these four phases (aggregates, cement paste, pores, and background) constitute four separate classes for image segmentation. However, the composition of the aggregate phase is not uniform, as demonstrated by the different intensity levels. Additionally, it is challenging to visually separate the aggregate and cement phases, especially when the aggregate size is small. Traditional segmentation methods include unsupervised clustering algorithms such as the Gaussian mixture model (GMM) or K-means, which are agnostic to spatial information, and therefore, will face challenges in segmenting the classes accurately. The complexity level of the present XCT data motivates the use of deep learning techniques to recognize image features and to label images accordingly.

In convolutional neural networks (CNNs), feature maps are created by running feature detectors or filters on the original images for edge detection and edge enhancement. The ensemble of feature maps constitutes a convolutional layer. Network training provides information about the most relevant features to search for in an image. When a feature detector is applied to an image, linearity in the filtered image may increase. Therefore, the nonlinearity of the image is enhanced by applying a rectified linear unit (ReLU) function to the feature maps. A pooling layer is then applied to ensure spatial invariance by downsampling each feature map according to a specified pooling operation, such as maximum pooling (max pooling) within a smaller filter window than the feature detector in the convolutional layer. By doing so, the pooling layer essentially extracts the dominant features of each feature map. A CNN typically includes a fully connected layer, which takes the flattened output of the last pooling layer in the network and applies an artificial neural network (ANN) to provide information on the likelihood of a specific class. Network training adjusts the probability of having a specific class with each iteration, typically using gradient descent. Several CNN architectures with different numbers of convolutional, pooling, and fully connected layers are available as described in the literature. A comprehensive introduction to CNN is provided in Wu [18]. CNNs are widely used for computer vision applications that are used for classification and object detection. Fully convolutional networks (FCNs) [19, 20] are a particular form of CNNs developed for semantic segmentation which do not have fully connected layers. Instead, these layers are replaced by 1×1 convolution layers to label each pixel (in 2D).

This section describes training of the FCN, U-Net [17], on few annotated slices. U-Net was then used to label all of the XCT images (1,018 slices in total) of a concrete specimen. This effort demonstrated that supervised learning provides significantly better results than conventional unsupervised segmentation. Traditional segmentation methods were used to provide an insight into the intensity classes of the original images by applying thresholding and anomaly detection. Less than 20 slices of the 3D XCT volume were labeled individually by attributing one of the five predetermined classes to each pixel in order to train a modified U-Net network on the features of the original images.

3.1 Application of Traditional Semantic Segmentation Methods

A first assessment of the segmentation effort was conducted [15] in which the intensity levels of grayscale XCT-based images were shown in a histogram. The low contrast between cement and aggregate phases

was shown in the form of a single intensity peak for both phases. An example original XCT-based image is shown in Figure 3. It was concluded that application of the clustering algorithm, DBSCAN combined with other image processing techniques, is the most promising approach outside of convolutional network-based image segmentation. The DBSCAN-based approach did not perform well near the edges of the specimen, where gray levels were slightly different, because DBSCAN considered the edge as cement. Other limitations are discussed in a previous report from Cheniour [15].

In this work, the same concrete specimen (Concrete-A [1]) was used to compare the results of the previous approach to those of GMM applied to intensity classes provided by multi-Otsu thresholding. GMM is an unsupervised clustering technique that fits multivariate Gaussians to data clusters—in this case pixels of similar intensity—and finds the corresponding Gaussian distribution parameters. GMM is typically used for anomaly detection. The multi-Otsu thresholding is an algorithm that can separate image pixels (intensity values). This approach is different from binary thresholding, which separates pixels into two intensity classes only.

Before any segmentation algorithm was applied, a block-matching, 3D filtering (BM3D) de-noiser was used to remove the noise and clean the data [21, 22]. Next, multi-Otsu thresholding of de-noised slices of the 3D sample were analyzed, showing that the data have at least four classes of gray levels, in addition to the background class. To that end, a combined GMM and multi-Otsu thresholding was applied to segment the given concrete XCT images. This resulted in results that are more consistent with features seen in the visual inspection of the images.

However, the GMM-segmented images had significant noise and misclassification, as seen in Figure 4a, which also results from the low contrast issue. The application of GMM on the DBSCAN output was also investigated, and the result is shown in Figure 4b. Performance of the latter approach was essentially still limited by the initial DBSCAN output. Therefore, standalone unsupervised clustering algorithms could not reach a high level of accuracy in the image segmentation given the challenging low contrast and wide range of cluster sizes.

3.2 Deep Learning-Based Semantic Segmentation

As demonstrated, traditional unsupervised algorithms failed to accurately segment the 3D XCT volumetric images of concrete samples. Considering the promising results obtained by deep neural networks for semantic segmentation in a range of scientific imaging applications [23–27], use convolutional neural networks were selected—specifically U-Net [17]—to perform semantic segmentation on these samples. The section below describes the U-Net architecture, the applied modification that was implemented to perform fast 3D segmentation with limited data, the network training procedure, and the results obtained by the trained network to date.

3.3 U-Net Architecture

U-Net is a widely used 2D semantic segmentation FCN that was first developed by [17] for the purpose of segmenting biomedical images. U-Net has a U-shaped architecture in which the left portion comprises a number of convolutions, Re-LU, and max pooling (contracting), and the right portion has convolutions, ReLU, and transposed convolutions (expanding). Finally, a 1×1 convolutional layer was applied such that






Figure 3. Example original XCT slice of the concrete specimen.



(a) GMM + multi-Otsu thresholding



(b) GMM + DBSCAN

Figure 4. Application of unsupervised clustering algorithms to segment the image shown in Figure 3.

each pixel is attributed a feature vector of the size of the number of classes. The total number of convolutional layers is 23 [17]. Since 2015, several U-Net-based architectures have been developed and applied as described in the literature, such as 3D U-Net for 3D semantic segmentation [28]. Example applications are provided in the literature [29–31]. Here, inspired by work performed by Ziabar [32], U-Net was modified to perform 3D semantic segmentation in a 2.5D scheme. The modified network architecture is shown in Figure 5. In this scenario, the input training data were divided into patches (subvolumes) of $N \times N \times 5$, with 5 corresponding to 5 neighboring slices. From each of the 5 input patch slices, the network learns to label the center patch slice. The key benefits of the proposed approach are as follows: (1) The network effectively captures 3D information without incurring any computational cost, because all of the intermediate convolutions remain the same as those in a 2D network, with only the first layer having 5 input channels instead of 1; and (2) the network is suitable for the limited 3D annotated data available for this effort, because large 3D CNNs require significant a amount of 3D labeled data; in this case, a limited number of labels are being used in the architecture. Furthermore, the loss function has been modified to calculate a weighted Dice loss [29] (see Eq. 1), as this is more suitable for imbalanced class labeling in this work.

$$WDC = \frac{1}{N} \frac{\sum_M (Y_{true} * Y_{pred} + \frac{\epsilon}{2})}{\sum_M (Y_{true} + Y_{pred} + \epsilon)}, \quad (1)$$

where N is the total batch size in each epoch, and Y_{true} and Y_{pred} are the one-hot encoded ground truth and predicted labels, respectively. Each voxel can belong to one of five classes. ϵ is a small constant that is used to avoid division by zero.




Figure 5. Schematic of the proposed 2.5D U-Net architecture.

3.4 Network Training

In order to train the 2.5D U-Net, some of the data needed to be annotated. Based on prior experience, it was determined that 50 slices of the 3D volume should be sufficient for the initial network training.

To annotate the data, the already weakly segmented data were leveraged by GMM+multi-Otsu thresholding. The steps taken for annotation are detailed here. First BM3D was applied to suppress noise in the volume. Then, the GMM+multi-Otsu threshold was applied to generate imperfect labels of the data. The weakly annotated data were then downsampled by $2\times$. An expert then manually corrected (pixel by pixel) the weak segmentation results in 24 slices of the downsampled data using the 3D Slicer software [33]. The 24 slices were chosen from the top (15 consecutive slices), bottom (3 consecutive) and center (6 consecutive slices) to account for intensity changes along the sample’s height. Once the labeling was complete, the labels back were linearly interpolated back to their actual resolution ($2\times$ up-sampling), which resulted in 48 labeled slices in total.

The four classes consisted of (1) dark areas (low intensity) typically surrounding pores, (2) the cement paste, (3) low-brightness aggregates, and (4) high-brightness aggregates (high intensity). Pores and background were easily detected by the thresholding method. The result of the manual segmentation of slice 210 (shown in Figure 3) is presented in Figure 6. The image shows the four classes (1–4), in addition to the background class (0). It is worth noting that manual segmentation is a very time consuming task and for this data, it took up to an hour per slice using 3D Slicer.

Finally, 3 additional slices were manually segmented to serve as a test data set following network evaluation. These slices were taken from different axial positions from the training slices to avoid bias in the network evaluation. The original XCT images toward the tops and bottoms of the samples were not used to train the model or evaluate its performance because they showed concentric circles of different intensities related to the instrument’s operation. These artifacts were not visible from slices 200 through 850.

The final training pairs were 48 slices of $1,004 \times 1,024$ from the XCT volume and their corresponding labels. A sliding window strategy was used to create 3D patches of training data. The patch size was 128×128 in the in-plane direction (X-Y) and 5 in the cross-plane (Z) direction. The stride for the patch size was 64 in the X-Y direction and 1 in the Z direction. Based on the analysis in [32], 5 neighboring slices were used. By using augmentations such as rotation (90, 180, 270) and image flipping, the training data were increased by $8\times$. Of these patches, 80% was used for training, and the remaining 20% was randomly selected for validation. The augmentation of the validation set was separated from training set, so no common data were included in the validation and training sets. To minimize the weighed Dice loss function, the Adam optimizer was used with an initial learning rate of 0.0001 and a default momentum setting. Validation loss was monitored, and if the validation loss stagnated for 15 consecutive epochs, then the learning rate was reduced by $2\times$. A batch size of 64 was used during training with an A100 machine with 2 GPUs. The choices of the initial learning rate, the batch size, and the patch window size was made after empirical study of the network and extensive training. In the future, use of a hyperparameter optimization approach [34, 35] could help to further improve the network.




Figure 6. Manual labeling of slice 210 shown in Figure 3.

3.5 Results and Performance Metrics

Upon network training and validation, the full volume (1,018 images of $1,004 \times 1,024$ pixels) was segmented slice by slice. Note that the model did not use the full images for training, as previously explained. The output images were slightly cropped after being segmented using the U-Net network, which is explained by the nature of U-NET [17]. The segmented images have a resolution of $992 \times 1,024$. Example segmented slices are shown in Figure 7. The dark intensity pixels surrounding pores (class 1) do not appear to be accurately detected. All labeled images from slice 200 to slice 850 were stacked as shown in Figure 8, in which class 0 is not shown to better visualize the reconstructed concrete volume. The figure shows a significant improvement in the identification of aggregate phases compared to DBSCAN's output in Cheniour's earlier report [15]. The aggregate edges are clearly defined, and the volume's edges show the presence of aggregates, whereas the DBSCAN-based segmentation failed to capture aggregates near the surface. The inner surface (or sample hole) primarily shows cement paste. This is likely caused by the difficulty in observing aggregates near the sample's hole during the manual segmentation process. Nonetheless, a visual inspection of the labeled volume shows excellent results compared to those obtained in the previous approach.

In the volume shown in Figure 8, the volume fractions of aggregates, cement, and pores are 61.7, 36.1, and 2.2%, respectively. These fractions demonstrate a notable improvement in the detection of small aggregates compared to the previously obtained 50% aggregate fraction [15]. It must be noted that small aggregates were quite challenging to visualize and label during the manual segmentation because their edges are not clear. This issue is essentially related to the original image rather than the segmentation method itself.

To quantify the performance of the U-Net model, the test data set was used, which consists of 3 additional manually labeled slices taken from different axial positions. The slice numbers are 342, 430, and 600. Using bilinear interpolation, the labeled images were up-sampled to match the original image's resolution,




Figure 7. Segmented XCT images using U-Net.




Figure 8. Clipped stack of labeled images (slices 200 to 850) using U-Net. Background voxels (class 0) were removed in this illustration for visualization purposes.

and then they were cropped to the dimensions of the U-Net output (992×1024). The resulting images were flattened into a concatenated 1D-array for performance assessment. This data array represents the ground truth (result of manual labeling). The U-Net outputs corresponding to the 3 test slices are also flattened and concatenated into a different 1D-array. This array represents the predicted segmentation. The *precision* (P), *recall* (R), and *F1-score* ($F1$) were measured first. These measures are defined such that

$$P = \frac{TP}{TP + FP}, \quad (2)$$

$$R = \frac{TP}{TP + FN}, \quad (3)$$

and

$$F1 = 2 \frac{P * R}{P + R}, \quad (4)$$

where TP is a true positive result (correctly identified as belonging to a specific class), FP is a false positive result (incorrectly identified as belonging to a specific class), and FN is a false negative result (incorrectly identified as not belonging to a specific class). Table 2 shows P , R , and $F1$ for each label class. Class 0 is the most accurately labeled one because it is the background, which is easily separable from the other classes. Classes 2 and 3 (cement and darker aggregates) are highly accurate. Class 4 has a particularly lower recall and F1-score. This could be explained by the different intensities within aggregates: that is, the brightness might differ even within a bright aggregate. The results show that class 1 is not detected properly by the network. This has two main reasons: (1) the very imbalanced annotation of

this class in comparison with other classes in the data set, and (2) the similarity between the regions labeled as class 1, with pores that are labeled as background, both in terms of shape and to some extent close gray level values. The accuracy is determined by the equation

$$accuracy = \frac{TP + TN}{TP + TN + FP + FN}, \quad (5)$$

where TN is a true negative results (correctly identified as not belonging to a specific class). The evaluation accuracy is 95.5% using the 3 additional slices. The macro average is taken as the arithmetic mean of the F1-scores of all classes. Because Class 1 was not detected, the macro average is only 73.3%. However, the weighted average is more relevant to this work because it accounts for the fraction of each class. The weighted average is 95.0%, which translates into a successful application of U-Net.

Table 2. Precision, recall, and F1-score evaluation

	Precision	Recall	F1-score
Class 0	0.975	0.999	0.987
Class 1	0.000	0.000	0.000
Class 2	0.904	0.947	0.925
Class 3	0.970	0.950	0.960
Class 4	0.926	0.697	0.795
Accuracy			0.955
Macro average	0.755	0.718	0.733
Weighted average	0.946	0.955	0.950

Finally, the intersection-over-union (IoU) score, which is commonly used to evaluate semantic segmentation, divides the number of intersections between the prediction and the ground truth over their number of unions. The IoU score in this case is 98.6%. Therefore, it is confidently concluded that the semantic segmentation using U-Net was successful. However, improvements must be made, especially to detect class 1, as discussed in Section 5. Overall, the training process took 76 seconds per epoch for 600 epochs. Testing on the full volume took about 30 seconds on the NVIDIA A100 system with 8 GPUs (during training only 2 GPUs were used).

3.6 Enhancing Accuracy Through Virtual Concrete Mesostructures

A major challenge to any AI-based approach is the large amount of training data needed to make reliable, conclusive decisions. It is the same for this effort of semantic segmentation. Samples were limited. For each sample, XCT scans were performed, and to ensure the best training results, a number of slices must be annotated manually (even if it is small). This process is tedious and time consuming. Annotation (manual or assisted by tools like DBSCAN) already includes some bias and is not 100% ground truth. Therefore, another modality of supportive data was developed as part of this work to train the AI segmentation system. This was done by using computationally generated concrete mesostructures that are faithful enough to be representative of real concrete. The main benefit of using these tools is that it makes a 100% ground truth data set available. This allows the architecture of the neural network to be adjusted (number of layers, their types, feedback loops, etc.) to achieve optimal segmentation results before real data are used. It also allows for the generation of a wide range of aggregate morphologies. For example, by controlling the generation

model parameters, elongated, disk-like, and round aggregates can be generated, as well as smooth or angular aggregates. Therefore, a rich and challenging training set can be produced, thus allowing the developed neural network to generalize beyond the limited experimentally available training set. Multiple computational tools exist as described in the literature that have been utilized in various concrete mechanics models [36–38]. Among these tools, the Concrete Mesostucture Generator (CMG) [2] is one of the most versatile open source packages available today in Python [39].

CMG uses the grain size distribution of the real concrete to create a virtual counterpart. Additional parameters enable CMG to introduce the various aggregate morphology characteristics. The most important parameters are the aspect ratios between the aggregate’s 3 major axes (envisioned as an ellipsoid), and the number of cuts or faces per aggregate. Varying these two parameters allows for representation of a wide range of aggregates as shown in Figure 9.

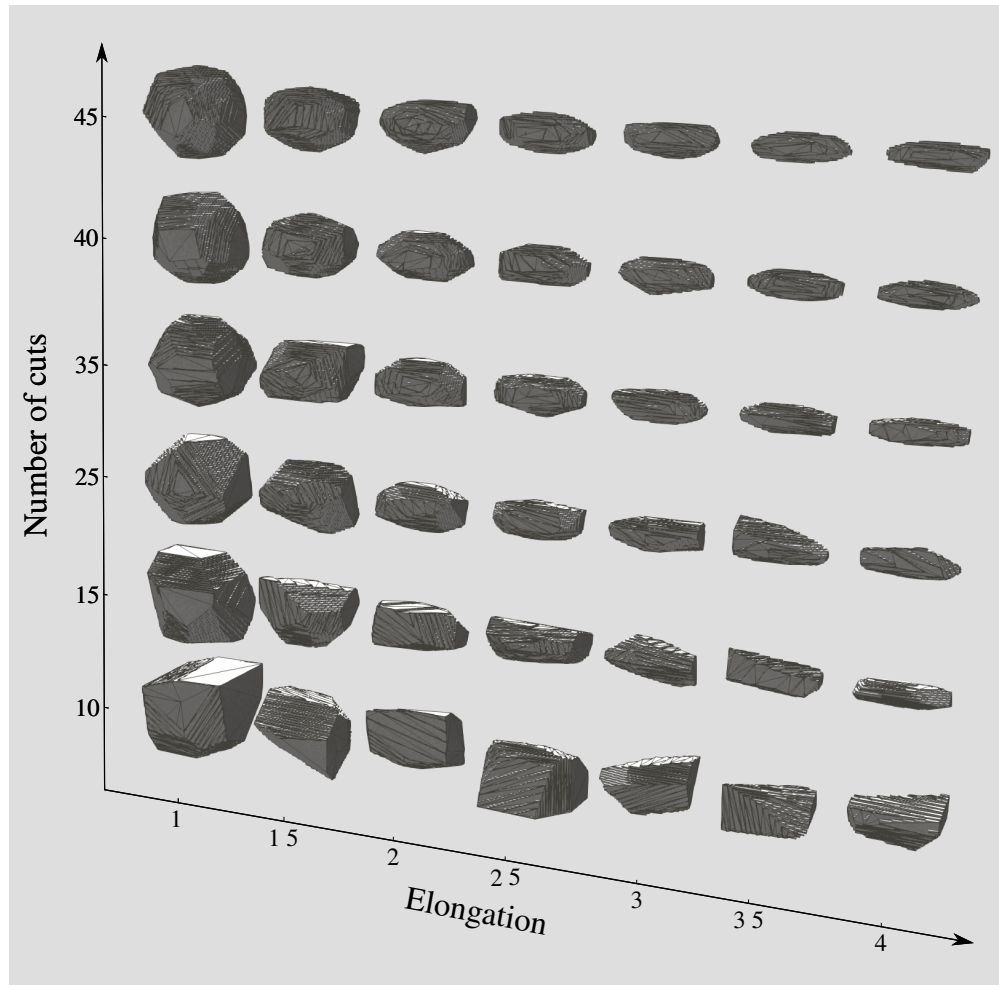


Figure 9. Influence of the number of cuts (N) and the elongation on the aggregate shape. (reproduced from [2]).

Other important parameters are the volume fraction of aggregate to be generated and the resolution of the generated 3D voxel system. It is important to train the segmentation procedure to homogenize fine

aggregates below a specific threshold into the matrix. Otherwise, aggregate interfaces can become blurred, or multiple aggregates can be connected because a small aggregate is filling in the space. This is accomplished manually during annotation by neglecting to annotate those aggregate pieces, using a smoothing filter, or both. The virtual mesostructure training set allows this feature to be directly built in by setting a lower threshold on the aggregate size, along with the correct aggregate volume fraction that would create a realistic interstitial space of mortar. The resolution is related to this because it must be finer than the minimum feature to be simulated. Figure 10 shows the combined effects of resolution and the number of cuts for cubic samples with volume fraction $V_f = 0.35$. The original code was modified to provide each aggregate piece with a unique ID so that particle colors can be varied based on their chemical compositions, making the mesostructure even more realistic.

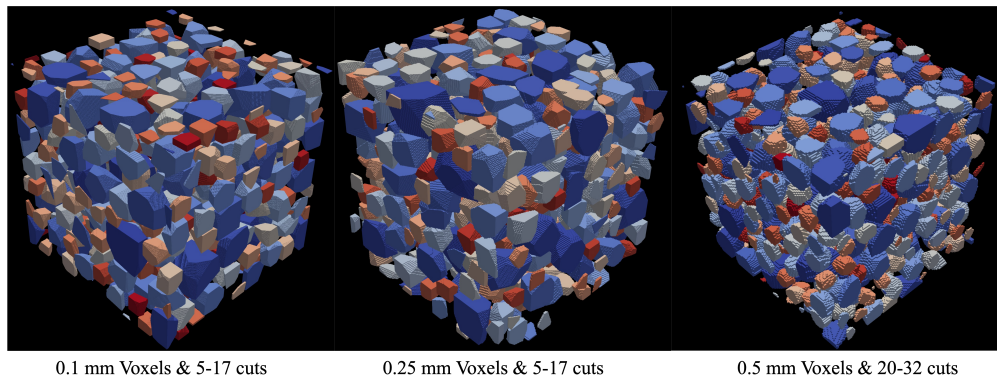


Figure 10. Influence of resolution and number of cuts on the aggregate shape.

To illustrate the combined effects of volume fraction and resolution with the same number of cuts, three different cases are shown in Figure 11.

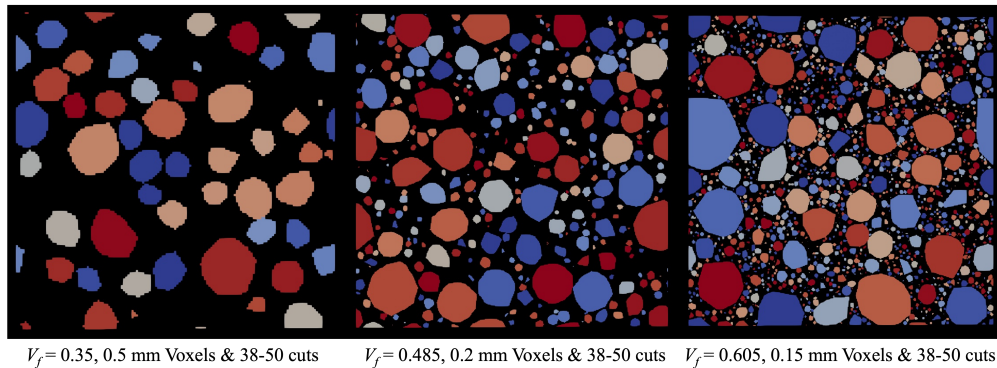


Figure 11. Influence of volume fraction (V_f) and resolution on the mesostructure.

Once a cubic sample is generated, simple extractions can be made to cut any shape. For example, Figure 12 shows a cylinder that is 40 mm in diameter and 60 mm in height virtually cored from a 60 mm cube. The coloring of the aggregate is accomplished by randomly assigning colors to the different aggregate IDs, showing the versatility in creating a wide range of mesostructure variations.




Figure 12. Coring a cylinder ($h = 60\text{mm}$, $\phi = 40\text{ mm}$) from a 60 mm cube with $V_f = 0.605$, 0.15 mm voxels and 38–50 cuts per aggregate.

4 APPLICATION TO IRRADIATED CONCRETE MODELING

This section describes use of the segmented XCT volume to serve as a 3D simulation domain in MOSAIC. Cheniour [15] details the approach to convert the stack of images to a 3D domain with grains included within aggregates. Essentially, the finite element-based Multiphysics Object-Oriented Simulation Environment (MOOSE) was used to produce a Voronoi diagram to account for particles within aggregates and interparticle interfaces (IPI). A Python script was then used to assign a specific phase or mineral to each particle according to the provided experimental aggregate composition. Finally, the phase data were converted to a MOSAIC-compatible phase map file and used as a simulation domain in the example provided in Cheniour [15]. In this work, the effort exclusively focused on converting the stack of images to a simulation domain without generating particles within aggregates. Given the high quartz content of Concrete-A aggregates [1], it was assumed that aggregates only contain quartz. These assumptions were made for the sake of utilizing the convolutional network results for 3D simulations in MOSAIC in a simplified and straightforward way. However, the methodology presented in Cheniour [15] is naturally applicable to the deep learning segmentation results and is recommended for the quantitative assessment of damage and expansion in irradiated concrete using 3D MOSAIC.

4.1 3D Simulation Setup

To benefit from the high-resolution labeled images obtained in Section 3, the top left quarter of each image was used to produce the simulation domain. The images were then slightly cropped to reduce the fraction of background. The resulting resolution was 489×477 pixels. Despite only considering a part of each image, the full stack of images of this resolution requires a substantial amount of memory. For instance, if the full specimen's volume is used, then the domain size becomes $489 \times 477 \times 1,018$, which incurs a significant computational cost. To reduce the memory requirement, the image resolution can simply be reduced as done in Cheniour [15]. For this work, the sample's height was limited to 50 slices, and the original resolution was maintained, as shown in Figure 13. Therefore, the domain size shrank to $489 \times 477 \times 50$, which enabled the use of a local ORNL Linux workstation with 64 GB of RAM. However, it would be beneficial to implement the use of the message passing interface (MPI) that is standard in MOSAIC to allow for the concurrent use of several compute nodes (or a computer cluster).

To generate a phase map for MOSAIC, a CSV file for each phase (quartz, cement paste, and pore/background) was obtained by reading the stack of images shown in Figure 13 and assigning each voxel the value of 1 if it corresponds to the specific phase, and 0 otherwise. All aggregate voxels (orange and red in Figure 13) were considered as quartz for simplicity. Pores and background were merged into the single phase shown in Figure 13 in blue. The remaining voxels corresponded to the hardened cement paste (hcp). The CSV files were then read by MOSAIC to generate an XML phase map, accounting for boundaries between aggregates and the cement paste. MOSAIC automatically creates a layer of interfacial transition zone (ITZ) that is one pixel thick, as shown in black in Figure 14.

The MOSAIC models and parameters were identical as the previous simulation [15] for the 3 phases included in this domain: quartz, hcp, and pores. In summary, an isotropic RIVE model was applied to the quartz phase based on the irradiation data contained in the Irradiated Minerals, Aggregates, and Concrete (IMAC) database [6]. The IMAC database also provides the elasticity tensors of a wide range of minerals, including quartz in this case. As in the previous milestone report, a linear viscoelasticity model was applied in the hcp and ITZ for simplicity. Note that MOSAIC also includes a logarithmic viscoelasticity model to




Figure 13. Reduced 3D simulation domain in MOSAIC (50 slices).




Figure 14. Phase map of the first layer of the 3D simulation domain.

account for long-term creep. The elastic properties of hcp and ITZ were also taken from the previous simulation in Cheniour’s earlier report [15]. For pores/background, a relatively low elastic modulus was assumed to ensure field continuity (stress and strain). Damage was not modeled to reduce the computational time of the simulation. Finally, the highest fluence was set to 0.15 n pm^{-2} at 36.3 days (accelerated irradiation conditions following data in [1]) and the temperature to 66.7°C . MOSAIC uses periodic boundary conditions. Free boundary conditions were used in this work.

4.2 Results

The total strain (%) in both the quartz and hcp phases is shown in Figure 15 for the xx [100], yy [010], and zz [001] directions. Because the RIVE model is isotropic, the total strain in all 3 directions is almost equivalent in quartz. In hcp, the strain is highest in the [001] direction. The stress (MPa) in the 3 directions is also plotted against fluence (n pm^{-2}) in Figure 16. The plots show that the stress is highest in the [001] direction and is comparable in the other directions considered here. Note that the height of the simulated domain is significantly smaller than the dimensions of the stacked images (almost one tenth of the image dimensions). It is crucial to include a damage model in future simulations because a previous 2D simulation of Concrete-A has shown that at 0.15 n pm^{-2} , the average damage d (calculated as the average of the IPI, ITZ, and hcp damage scalars) reaches $d \approx 0.44$. The simulation time was 11h:31m:23s with 8 threads. The numerical convergence greatly improved compared to Cheniour’s previous simulation [15].

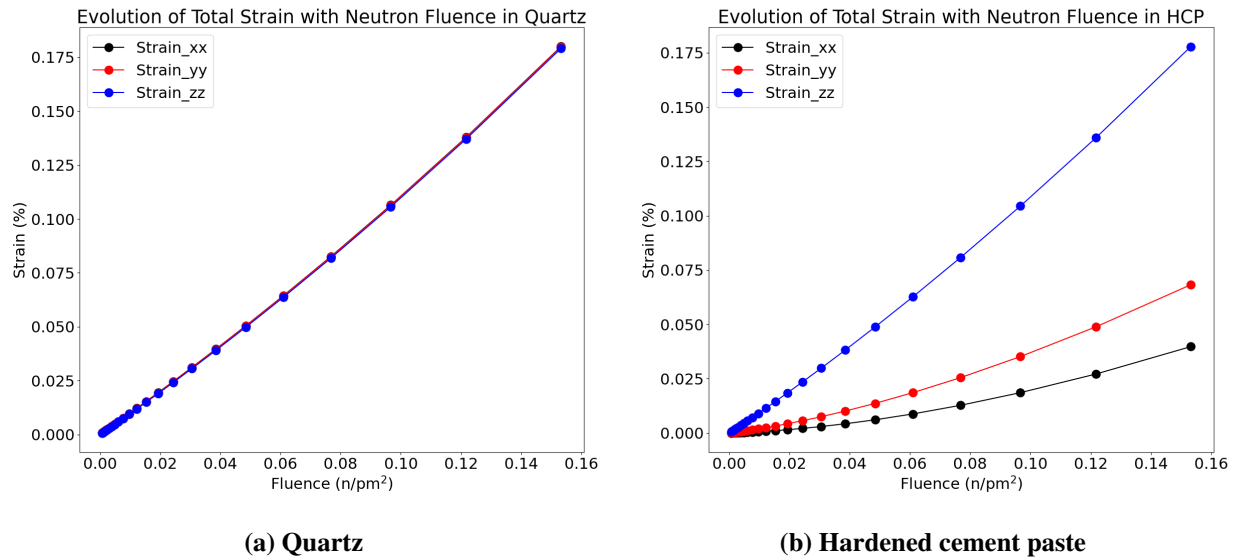

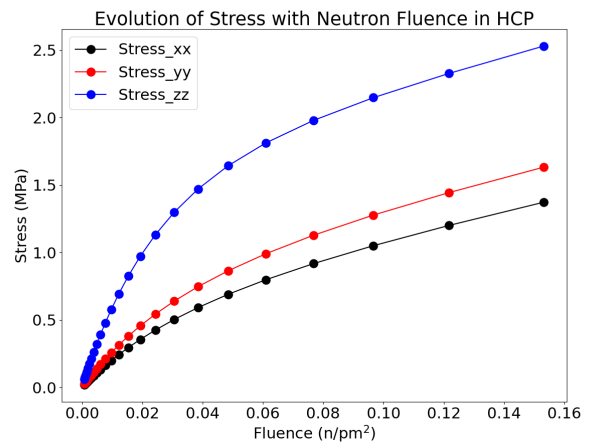


Figure 15. Evolution of strain (%) in the xx, yy, and zz directions with fluence (n pm^{-2}) in quartz and hcp as predicted by MOSAIC.



(a) Quartz



(b) Hardened cement paste

Figure 16. Evolution of stress (MPa) in the xx, yy, and zz directions with fluence (n pm⁻²) in quartz and hcp as predicted by MOSAIC.

5 DISCUSSION

The XCT-generated images clearly separate the voids and the surrounding area of the concrete specimen; however, it is difficult to distinguish the aggregates from the cement paste. The previous milestone report [15] employed a clustering algorithm (DBSCAN) to detect aggregates and label the phases present in each image. However, several limitations were reported, and convolutional networks were considered as an alternative image segmentation method. In this report, a modified version of the U-Net convolutional network architecture was used. The U-Net-based network was used to perform image segmentation to label the XCT-based concrete images. The segmentation results showed great accuracy (95.5%) when compared to manually labeled images. The accuracy of the concrete microstructure representation can still significantly benefit from additional data such as NCT, which is scheduled to occur early August 2022. In fact, NCT is sensitive to H atoms, which results in an improved contrast between aggregates and the cement paste. This is significant, because during the manual labeling of images (to generate a training data set), it was sometimes challenging to see boundaries between small aggregates with hcp. The combination of XCT and NCT may enable the use of traditional and simple segmentation techniques such as thresholding or clustering algorithms to some extent.

Although the present approach constitutes a milestone in the accurate segmentation of XCT-based concrete images, it is important to note that this method is time consuming because the generation of the network training data set requires manual labeling of several high-resolution images. An alternative approach would be to generate synthetic concrete microstructures that resemble the real microstructures enough to be used as a training data set for the segmentation network. This is possible because of the availability of data on aggregate morphology in simulators of concrete mesostructures such as PyCMG [39].

In addition to the challenge of segmenting XCT images, running 3D simulations of irradiated concrete in MOSAIC is not straightforward and still requires code development to include MPI capabilities and to use computer clusters. In fact, the simulation presented in Section 4 required nearly 73% of the 64 GB available memory, despite the substantial reduction in the domain size and the absence of a damage model. This example simulation demonstrates the need to extend the parallelization capabilities to utilize several machines simultaneously, which would increase the available memory. The current version of MOSAIC uses OpenMP, which is an application programming interface that supports shared-memory applications (single compute node use).

Eventually, it is anticipated that 3D simulations of irradiated concrete in MOSAIC will significantly improve the predictions of damage initiation and propagation in a concrete specimen. A previously developed methodology accounts for the presence of small particles within aggregates, as well as aggregate compositional data in the reconstructed 3D microstructures [15]. It is envisaged that, upon further code development for the parallelization of MOSAIC, 3D FFT-based simulations using the newly generated realistic simulation domain will be run on computer clusters and documented in future publications.

6 CONCLUSION

This report provides an improved methodology to utilize XCT-based concrete images in the generation of 3D domains for irradiated concrete simulations in MOSAIC. The proposed methodology uses a modified version of the convolutional network architecture U-Net to label the XCT images. These images show low contrast between aggregates and the cement paste. The segmentation results showed an evaluation accuracy of 95.5%. The labeled images were then stacked to generate the 3D high-resolution simulation domain. A simple example simulation of irradiated concrete using a portion of the generated domain was provided in this report. It is imperative to extend MOSAIC parallel capabilities through further code development to benefit from ORNL's high performance computing resources such as computer clusters. This advancement would enable simulation of large high-resolution 3D concrete microstructures. This report also includes recommendations on alternative approaches to overcome the time constraint of manual segmentation. Alternatives include the combination of XCT with NCT data, which may allow for the use of traditional unsupervised segmentation methods (clustering, thresholding), as well as the generation of synthetic microstructures using available data on aggregate morphology to use as a training data set for the convolutional network. NCT is scheduled for early August 2022.

7 REFERENCES

References

- [1] I. Maruyama, O. Kontani, M. Takizawa, S. Sawada, S. Ishikawa, J. Yasukouchi, O. Sato, J. Etoh, and T. Igari. Development of the Soundness Assessment Procedure for Concrete Members Affected by Neutron and Gamma-Irradiation. *Journal of Advanced Concrete Technology*, 15:440–523, 2017.
- [2] Vijaya Holla, Giao Vu, Jithender J. Timothy, Fabian Diewald, Christoph Gehlen, and Günther Meschke. Computational Generation of Virtual Concrete Mesostructures. *Materials*, 14(14):1–19, 2021.
- [3] H. Graves, Y. Le Pape, D. Naus, J. Rashid, V. Saouma, A. Sheikh, and J. Wall. Expanded Material Degradation Assessment (EMDA), Volume 4: Aging of Concrete. Technical Report NUREG/CR-7153, ORNL/TM-2011/545, US Nuclear Regulatory Commission, 2014.
- [4] P. Stanton. Expansion of Concrete through Reaction between Cement and Aggregate. *Proceedings of the ASCE*, 66:1781–1811, 1940.
- [5] A. Abd-Elssamd, Z.J. Ma, H. Hou, and Y. Le Pape. Influence of Mineralogical and Chemical Compositions on Alkali-Silica Reaction of Tennessee Limestones. 261:119916.
- [6] Y. Le Pape, M.H.F. Alsaid, and A.B. Giorla. Rock-Forming Minerals Radiation-Induced Volumetric Expansion—Revisiting the Literature Data. *Journal of Advanced Concrete Technology*, 16:191–209, 2018.
- [7] Y. Le Pape, J. Sanahuja, and M.H.F. Alsaid. Irradiation-Induced Damage in Concrete-Forming Aggregates—Revisiting Literature Data through Micromechanics. *Materials and Structures*, 53:35, 2020.
- [8] Yuliia Khmurovska and Petr Štemberk. RBSM-Based Model for Prediction of Radiation-Induced Volumetric Expansion of Concrete Aggregates. *Construction and Building Materials*, 294:123553, 2021.
- [9] Y. Le Pape, K.G. Field, and I. Remec. Radiation Effects in Concrete for Nuclear Power Plants – Part II: Perspective from Micromechanical Modeling. *Nuclear Engineering and Design*, 282:144–157, 2015.
- [10] A. Giorla, Y. Le Pape, and C. Dunant. Computing Creep-Damage Interactions in Irradiated Concrete. *Journal of Nanomechanics and Micromechanics*, 7(2), 2017.
- [11] F. Chen, J. Sanahuja, B. Bary, and Y. Le Pape. Effects of Internal Swelling on Residual Elasticity of a Quasi-Brittle Material through a Composite Sphere Model. *International Journal of Mechanical Sciences*, page 107390, 2022.
- [12] Fuqiang Gao, Doug Stead, and Davide Elmo. Numerical Simulation of Microstructure of Brittle Rock Using a Grain-Breakable Distinct Element Grain-Based Model. *Computers and Geotechnics*, 78:203–217, 2016.

- [13] C.E. Torrence, A.B. Giorla, Y. Li, E. Tajuelo-Rodriguez, J.D. Arregui-Mena, T.M. Rosseel, and Y. Le Pape. MOSAIC: An Effective FFT-based Numerical Method to Assess Aging Properties of Concrete. *Journal of Advanced Concrete Technology*, 19(2):149–167, 2021.
- [14] Yahui Zhang and Louis Ngai Yuen Wong. A Review of Numerical Techniques Approaching Microstructures of Crystalline Rocks. *Computers and Geosciences*, 115(March):167–187, 2018.
- [15] Amani Cheniour, Yujie Li, Yann M. Le Pape, Ercan Cakmak, Jose David Arregui-Mena, Elena Tajuelo Rodriguez, and Thomas M. Rosseel. Development of a Reconstruction Methodology Based on X-Ray Computed Tomography to Generate Realistic 3D Concrete Microstructures in MOSAIC. Technical Report ORNL/TM-2021/2156, Oak Ridge National Laboratory, 8 2021.
- [16] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *kdd*, volume 96, pages 226–231, 1996.
- [17] O. Ronneberger, P.Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]).
- [18] Jianxin Wu. Introduction to Convolutional Neural Networks. *National Key Lab for Novel Software Technology. Nanjing University. China*, 5(23):495, 2017.
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [20] Juntang Zhuang, Junlin Yang, Lin Gu, and Nicha Dvornek. Shelfnet for Fast Semantic Segmentation. *Proceedings - 2019 International Conference on Computer Vision Workshop, ICCVW 2019*, pages 847–856, 2019.
- [21] Kostadin Dabov, Alessandro Foi, and Karen Egiazarian. Video Denoising by Sparse 3D Transform-Domain Collaborative Filtering. *European Signal Processing Conference*, 16(8):145–149, 2007.
- [22] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. BM3D Image Denoising with Shape-Adaptive Principal Component Analysis. In *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- [23] Graham Roberts, Simon Y Haile, Rajat Sainju, Danny J Edwards, Brian Hutchinson, and Yuanyuan Zhu. Deep learning for semantic segmentation of defects in advanced stem images of steels. *Scientific reports*, 9(1):1–12, 2019.
- [24] Renata Lorenzoni, Iurie Curosu, Sidnei Paciornik, Viktor Mechtcherine, Martin Opperma, and Flavio Silva. Segmentation of the Micro-Structure of Strain-Hardening Cement-Based Composites (SHCC) by Applying Deep Learning on Micro-Computed Tomography Scans. *Cement and Concrete Composites*, 108:103551, 2020.
- [25] Luke Scime, Derek Siddel, Seth Baird, and Vincent Paquit. Layer-Wise Anomaly Detection and Classification for Powder Bed Additive Manufacturing Processes: A Machine-Agnostic Algorithm for Real-Time Pixel-Wise Semantic Segmentation. *Additive Manufacturing*, 36:101453, 2020.

- [26] Yu Song, Zilong Huang, Chuanyue Shen, Humphrey Shi, and David A Lange. Deep Learning-Based Automated Image Segmentation for Concrete Petrographic Analysis. *Cement and Concrete Research*, 135:106118, 2020.
- [27] Alberto Garcia-Garcia, Sergio Orts-Escolano, Sergiu Oprea, Victor Villena-Martinez, and Jose Garcia-Rodriguez. A Review on Deep Learning Techniques Applied to Semantic Segmentation. *arXiv preprint arXiv:1704.06857*, 2017.
- [28] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [29] Amirkoushyar Ziabari, Abbas Shirinifard, Matthew R. Eicholtz, David J. Solecki, and Derek C. Rose. A Two-Tier Convolutional Neural Network for Combined Detection and Segmentation in Biological Imagery. *GlobalSIP 2019 - 7th IEEE Global Conference on Signal and Information Processing, Proceedings*, pages 1–5, 2019.
- [30] Xiaomeng Li, Hao Chen, Xiaojuan Qi, Qi Dou, Chi-Wing Fu, and Pheng-Ann Heng. H-DenseUNet: Hybrid Densely Connected UNet for Liver and Tumor Segmentation from CT Volumes. *IEEE transactions on medical imaging*, 37(12):2663–2674, 2018.
- [31] Libin Jiao, Lianzhi Huo, Changmiao Hu, and Ping Tang. Refined UNet: UNet-Based Refinement Network for Cloud and Shadow Precise Segmentation. *Remote Sensing*, 12(12):2001, 2020.
- [32] Amirkoushyar Ziabari, Dong Hye Ye, Lin Fu, Somesh Srivastava, Ken D Sauer, Jean-Baptiste Thibault, and Charles A. Bouman. Model Based Iterative Reconstruction with Spatially Adaptive Sinogram Weights for Wide-Cone Cardiac CT. In *Proc. 5th Intl. Mtg. on Image Formation in X-ray CT*, pages 15–19, Salt Lake City, Utah, 2018.
- [33] Ron Kikinis, Steve D. Pieper, and Kirby G. Vosburgh. *3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support*, pages 277–289. Springer New York, New York, NY, 2014.
- [34] Maryam Parsa, Aayush Ankit, Amirkoushyar Ziabari, and Kaushik Roy. PABO: Pseudo Agent-Based Multi-Objective Bayesian Hyperparameter Optimization for Efficient Neural Accelerator Design. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–8, 2019.
- [35] Maryam Parsa, John P. Mitchell, Catherine D. Schuman, Robert M. Patton, Thomas E. Potok, and Kaushik Roy. Bayesian Multi-objective Hyperparameter Optimization for Accurate, Fast, and Efficient Neural Network Accelerator Design. *Frontiers in Neuroscience*, 14, 2020.
- [36] Huan He, Piet Stroeven, Eric Pirard, and Luc Courard. On the Shape Simulation of Aggregate and Cement Particles in a DEM System. *Advances in Materials Science and Engineering*, 2015, 2015.
- [37] Jing Hu and Piet Stroeven. Shape Characterization of Concrete Aggregate. *Image Analysis and Stereology*, 25(1):43–53, 2006.
- [38] Numerical Simulation-Based Damage Identification in Concrete. *Modelling*, 2(3):355–369, 2021.
- [39] Concrete Mesostructure Generation Using Python - PyCMG 1.0 Documentation.