Helium Bubble Evolution in Nano-Infiltrated Transient Eutectic SiC Utilizing Machine Learning Analysis

A Thesis Presented for the

Master of Science

Degree

The University of Tennessee, Knoxville

Riley “Kip” Jean Wheeler

May 2025

For those who believed in me. Thank you to my sunshine as well as my mother and grandmother.

You brought me into the world for a reason and it was science.

Acknowledgements

Special thanks to Elijah Davis and Steven Frankowski, and the rest of the Tennessee Ion Beam Laboratory team for contributing assistance and training for the experimental work, specifically the XRD analysis and nanoindentation, as well as continuous support throughout my research. Special thanks to Christopher Field from Theia Scientific LLC for contributing assistance and training with machine learning models. A lot of my initial model work can be attributed to Dr. Field's expertise and a lot of opportunities have come from our work together. Special thanks to Caen Ang from the High Temperature Materials Laboratory at UT for contributing samples and valued expertise on processing and experimentation. Special thanks to Yongqiang Wang from the Ion Beam Materials Laboratory at Los Alamos National Laboratory for the great work done for the experimental implantation. Special thanks to Eric Lang of the University of New Mexico and his cohorts (Nathan Madden, Ryan Schoell, Ryan Peña, David J. Señor, and Andrew Casella) for contributing the lithium aluminate data for the preliminary models. Special thanks to Kory Burns of the University of Virginia and cohorts for the palladium nickel data used for the preliminary models. Special thanks to Michael Koehler from the UT Institute for Advanced Materials and Manufacturing for assistance with the XRD machinery and analysis. Finally, thank you so much to my committee, consisting of Sergei Kalinin, Xingang Zhao, Christopher Field, and Khalid Hattar, for their assistance and the confidence to finish this work.

This work was performed, in part, at the Center for Integrated Nanotechnologies, an Office of Science User Facility operated for the U.S. Department of Energy (DOE) Office of Science. Los Alamos National Laboratory, an affirmative action equal opportunity employer, is

managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA, under contract 89233218CNA000001.

Abstract

The nuclear energy industry is searching for materials that can be used in new reactor designs, such as fusion energy systems, and withstand the volatile conditions. Silicon carbide (SiC) has emerged as a versatile material that is already used in fission reactors and could be applied to the fusion reactor design due to its significant strength, thermal stability, and resistance to radiation damage. Characterization of materials results in massive amounts of data due to the ever-improving microscopy technology. Micrographs taken of samples implanted with helium bubbles can have thousands of microscopic features that need to be annotated. This large amount of data and human error are a limiting factor on data analysis in the nuclear materials science field. Many machine learning (ML) model approaches have been explored to assist human researchers process visual data faster than manual annotation. Models such as You Only Look Once (YOLO) have been used to identify features such as black dots, cavities, and grains. With the purpose of addressing the growing need for radiation stable materials, a comparison of two production methods of SiC, chemical vapor deposition (CVD) and nano-infiltration transient eutectic (NITE), under helium implantation at three different fluences (1×10^{14}, 1×10^{15}, and 1×10^{16} ions/cm^2) was conducted. In addition, a ML model was made to evaluate the bubble size and density of the data. Nanoindentation was also performed to evaluate the mechanical stability of the materials. The insight gained into microstructural and nanomechanical properties during various helium irradiations helps to elucidate the structure-property relationship of NITE SiC during potential future fusion heat blanket applications.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

As the global demand for energy continues to rise, the need for clean, sustainable, and high-output energy sources has never been more critical. While renewables such as wind and solar play an important role in decarbonizing the energy sector, their reliance on environmental conditions make them challenging to implement as standalone solutions. Nuclear energy, with its unmatched energy density and ability to operate continuously, presents itself as a vital contributor to the future energy mix. In particular, nuclear fusion, the process that fuels stars, has long been recognized for its potential to provide virtually limitless energy with minimal environmental impact.

However, these advantages come with steep technological challenges. The extreme environment within a fusion reactor places immense demands on the materials used, particularly in components that come into direct contact with the plasma. To withstand these conditions, materials must exhibit exceptional thermal stability, radiation resistance, and structural integrity. Among the materials under consideration, silicon carbide (SiC) stands out due to its high melting point, excellent thermal conductivity, low neutron activation, and strong mechanical properties. SiC is especially promising for use in blanket systems, which serve to capture heat, breed tritium, and shield other reactor components from neutron damage. In silicon carbide's composite form (SiC/SiC), the material gains additional toughness and reliability, making it a strong candidate for fusion energy systems.

In addition to thermal and structural challenges, fusion environments expose materials to high-energy neutrons that generate helium within the lattice. Helium atoms can migrate, coalesce, and form pressurized bubbles that weaken the material by causing swelling, embrittlement, or fracture. Understanding and characterizing helium bubble formation is

essential for predicting the long-term behavior of structural components and ensuring the safety and longevity of future fusion reactors. Advanced microscopy techniques, including transmission electron microscopy (TEM) and scanning TEM (STEM), are employed to observe these microstructural changes at the nanoscale. However, as instruments become more powerful, they produce overwhelming volumes of data, creating a bottleneck in analysis. To address this, researchers are increasingly turning to machine learning (ML). Supervised algorithms and neural networks are now being trained to automatically detect features such as voids, cracks, or helium bubbles within micrographs. This is done in an effort to significantly reducing the time required for characterization and improving reproducibility.

By combining the promise of fusion energy with next-generation materials like SiC and leveraging tools such as machine learning for microstructural analysis, the field of nuclear materials science is entering a transformative era. These innovations not only bring the vision of clean fusion energy closer to reality but also open the door to broader advancements in high-performance materials engineering.

Chapter 2: Background

Silicon Carbide for Advanced Nuclear Applications

The world and its expanding technology bring with them a growing demand for energy and growing concerns about carbon emissions and their effects on the environment. This has led to a greater focus on sustainable and reliable energy sources, such as nuclear energy. Nuclear energy has long been considered a sustainable energy source with its high energy density and near-zero carbon emissions when in operation. Many organizations, such as the Massachusetts Institute of Technology (MIT), have long been releasing reports about the pivotal role nuclear technologies could play in meeting future energy needs and environmental concerns [1]. Their 2011 report emphasizes the importance of improvements in nuclear fuel cycles, reactor safety, and waste management to contribute to a sustainable energy future. The most widely used nuclear power plants are pressurized water reactors (PWRs), which utilize Uranium-235 as a fuel for fission reactions that generate heat, which is used to turn turbines and produce electricity. According to the International Atomic Energy Agency (IAEA), nuclear power already contributes 10% of the world's total electricity and nearly 30% of all low-carbon electricity generation [2]. **Figure 1** below shows the statistics on most forms of energy production to compare the amount of carbon dioxide per equivalent kilowatt hour and reveals that nuclear is second only to river based hydropower in lowest emissions. This reaffirms nuclear fission the dual role of power in stabilizing energy production as well as mitigating carbon effects on the environment. Despite the clear advantages, nuclear energy is not without its challenges. Public fears, waste disposal, safety concerns, and high initial capital requirements have inhibited the wider adoption of nuclear-based power plants. The nuclear power industry strives to bring

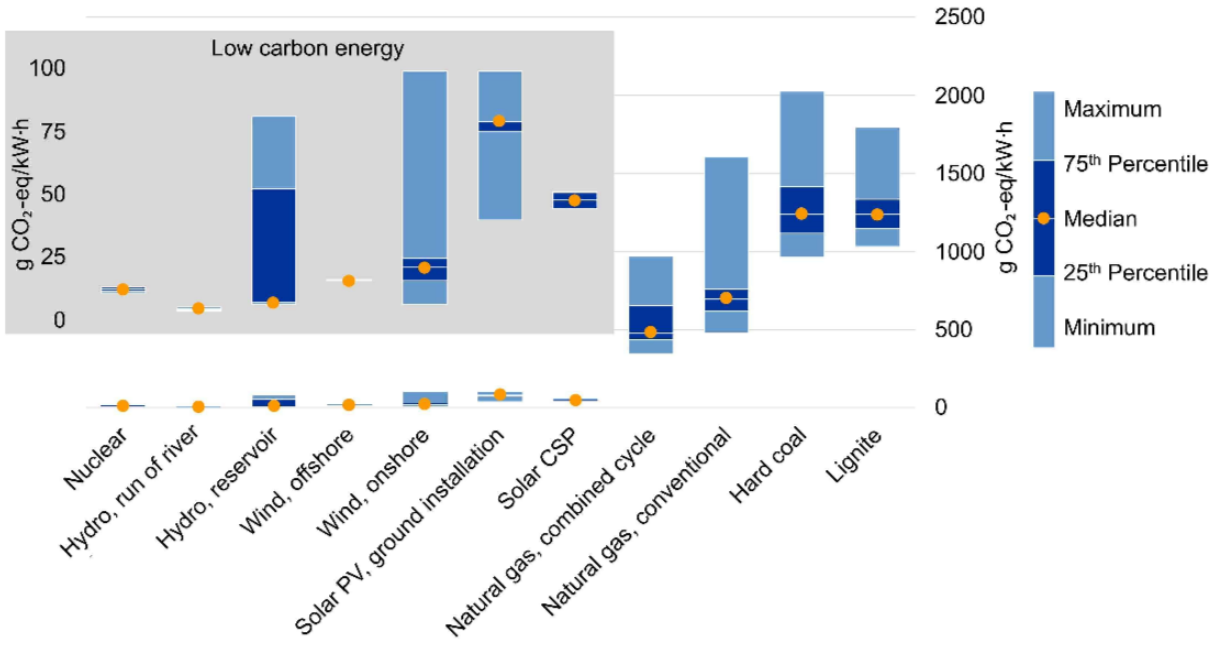

FIGURE 1. GREENHOUSE GAS EMISSIONS OF DIFFERENT ELECTRICITY GENERATION TECHNOLOGIES [2].

advancements in reactor technology and improve regulatory frameworks to produce ever safer and more efficient designs. One promising area of innovation is the development of nuclear fusion reactors, which have the potential to provide even greater energy production with less waste than traditional fission-based systems.

Nuclear fusion is a similar concept to fission, but instead of a heavier nucleus breaking apart and releasing energy, fusion sees two light nuclei combining to form a heavier nucleus. This holds immense promise as an energy source because the reaction generates more energy than fission while producing minimal long-lived radioactive waste, which is one of the biggest concerns involving fission. Fusion reactions release tremendous amounts of energy by fusing isotopes of hydrogen, such as deuterium and tritium. The primary byproduct of this reaction is helium, a harmless and non-toxic gas [3]. Taking these factors into account, the potential advantages of fusion over fission are significant. Fusion fuel is abundant and widely available as

deuterium contains a natural isotope of hydrogen that can be extracted from water, while tritium can be bred from lithium-6. Additionally, fusion reactions do not produce high-level waste (HLW). This makes long-term storage and disposal less problematic [3]. If utilized correctly, fusion could provide a virtually limitless supply of clean energy and revolutionize global energy systems while actively combating and aiding in climate change mollification efforts. Despite this, nuclear fusion remains a highly challenging endeavor due to the extreme conditions for sustained fusion. The extremely high temperatures (over 100 million degrees Celsius) and pressure require advanced materials and technologies. One of the most critical components of a fusion reactor is the choice of materials used to contain and manage the extreme conditions within the reactor core. This is most true in the case of Generation IV reactors and fusion reactors whose materials have yet to be chosen due to difficulties associated with simulating a true reactor environment. The Bumpy Tarus, as shown below, and later ELMO concepts were utilized to establish magnetic plasma confinement techniques. The Tokamak model was later selected as a more viable design, but the original Bumpy Taurus prototype sits outside of the Zeanah Engineering Complex to this day and can be seen below in **Figure 2**.

The development of fusion reactors faces numerous material science challenges as the reactor components must withstand intense heat, high neutron fluxes, and prolonged exposure to radiation without degrading. These conditions can cause structural damage, swelling, and embrittlement in conventional materials, compromising reactor safety and efficiency. As a result, the selection of materials revolves around finding economic containment materials that have exceptional thermal, mechanical, and radiation-resistant properties. Key components of a fusion reactor include the plasma-facing walls, structural supports, and the blanket system designed to capture and utilize the heat generated by fusion reactions. Materials used in these components

FIGURE 2. NASA'S LEWIS BUMPY TORUS USED TO TEST MAGNETIC PLASMA CONFINEMENT FOR FUSION SYSTEMS.

must have high thermal conductivity, low neutron activation, and resistance to radiation-induced defects to avoid degradation of the system [4]. The high start-up capital required for nuclear fission reactors tends to result in plans that are well understood, long-lasting, and inexpensive. Stainless steel is a major component in all modern reactors because of these very reasons. For the fusion reactor, silicon carbide (SiC) has emerged as a leading candidate material due to its unique combination of properties. **Figure 3** below shows a comparison of the dpa and temperatures associated with current fission reactors as well as fusion reactor concepts.

SiC is a ceramic material known for its strength, thermal stability, and resistance to radiation damage which makes it an ideal candidate for use in fusion reactor environments. SiC has a high melting point (2,700 ^{o}C), high thermal conductivity, and a low neutron absorption cross-section, all of which contribute to its ability to perform under extreme conditions, such as fusion reactor settings [5]. One of the primary roles of SiC in fusion reactors is within the heat blanket system. The heat blanket serves multiple purposes: it captures the heat generated by fusion reactions, breeds tritium through interactions with lithium, and shields other reactor components from neutron damage. SiC can be used in the form of fiber-reinforced composites

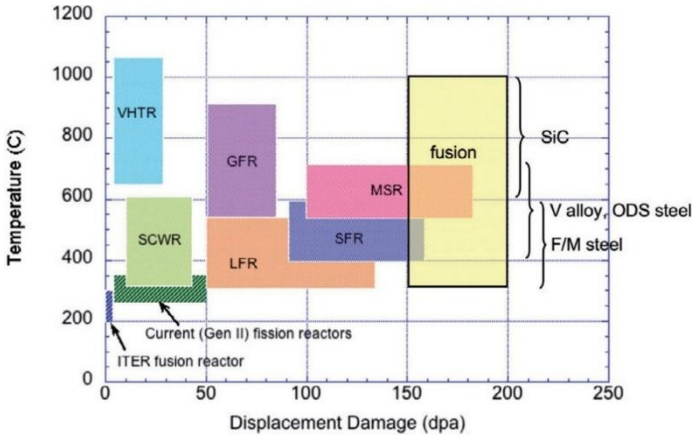

FIGURE 3. OVERVIEW OF OPERATION TEMPERATURES AND DISPLACEMENT DOSE REGIMES FOR MATERIALS IN CURRENT AND FUTURE FISSION AND FUSION ENERGY SYSTEMS [4].

(SiC/SiC), which offer enhanced mechanical properties and structural integrity compared to monolithic ceramics [5]. In addition to its thermal and mechanical properties, SiC exhibits excellent radiation tolerance. Research has shown that SiC can maintain its structural integrity even after prolonged exposure to high neutron fluxes, making it a reliable material for long-term fusion reactor operation [6]. The Very High Temperature Reactor (VHTR) utilizes SiC in its fuel and was able to prove reliability in conditions up to 1000 ^{o}C and around 20 dpa in the core [7-9]. Furthermore, the low activation properties associated with SiC reduce the risk of generating secondary radioactive waste. The SiC components in reactors have to be specially manufactured into their intended shapes which has led to a lot of different production methods of SiC that try to maximize density and purity while minimizing the cost associated with production. **Figure 4** below shows a schematic from 1985 of a fusion reactor concept utilizing SiC as the body of the reactor to act as a thermal blanket and facilitate better cooling of the core.

The most common way to produce SiC components is chemical vapor deposition (CVD) which involves using gas phase reactants onto hot bulk surfaces to form desired phases of SiC

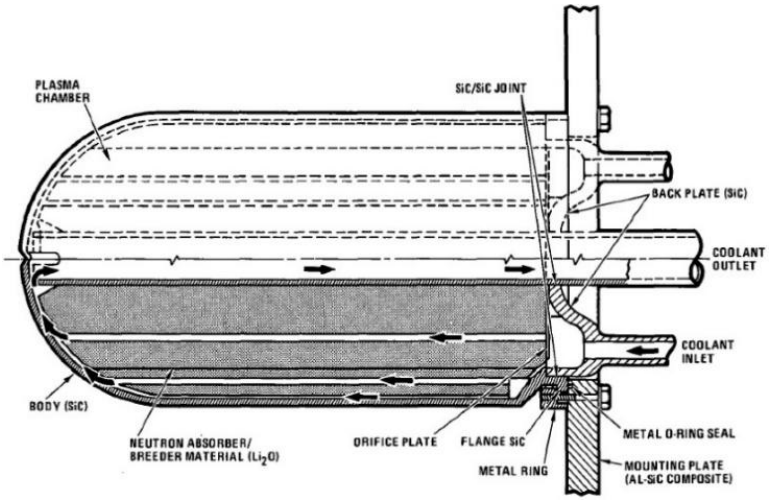

FIGURE 4. DUAL-COOLED FIRST WALL/BLANKET MODULE CONCEPT WITH SIC COMPOSITE STRUCTURE [5].

[10]. This process is very similar to chemical vapor infiltration (CVI) which also involves using gas phase reactants; however, with infiltration of Polya porous medium is used to form the desired SiC phase. Due to the immense similarity in production and resulting SiC, CVD and CVI will be used interchangeably in this background though the samples used in this work were CVD. The positives of this method are that it is very well established and that it produces very pure, predictable samples. The negative aspects of CVI revolve around the high cost, extended production time of weeks to months, and limited hermeticity under stress [11]. Polymer infiltration and pyrolysis (PIP) is the process in which specialized fibers in resin are immersed into preceramic polymer which is then heated in an anaerobic environment until the preceramic hardens into ceramic. The polymer infiltrations and pyrolysis has to be repeated several times to achieve the required density. The advantages associated with PIP are relatively low temperature production which prevents fiber damage, produces high purity results, and it can form a wide variety of matrices. The disadvantages associated with PIP revolve around the high cost and long production time as well as residual porosity which weakens the structure of the SiC [11, 12]. Reactive melt infiltration (RMI) involves infiltrating a porous reinforcing preformed with liquid metal. A relevant example to this study is the liquid silicon infiltration (LSI) to insert silicon into a carbon matrix and produce a fine grained SiC structure. The advantages of LSI are that is has a low cost, short production time and results in a product with high thermal and electrical conductivity. The shapes can also be complex when this method is used. The disadvantages of LSI are that the higher temperature can damage fibers, residual free silicon is present in the matrix, and the mechanical properties such as strength and modulus of elasticity were lower than typical SiC [11, 13]. NITE is yet another method of SiC production aimed at minimizing porosity and maximizing strength and thermal conductivity. **Figure 5** below

demonstrates the CVD technique (5a) as well as the NITE technique (5b) equipment set up to give an idea of the relative complexity of each process.

NITE SiC was developed to preserve the required performance factors while reducing porosity, which is left in similar methods such as CVI. Minimizing the porosity of SiC is necessary as it determines the maximum applicable stress and thermal conductivity [14]. NITE SiC has a small amount of additives which are able to crystallize during liquid phase sintering. It has been reported that because of these crystalline oxides, high temperature creep in NITE SiC is controlled by lattice diffusion and climb-controlled dislocation guide instead of the typical grain boundary sliding [15]. The thermo-structural aspects of NITE SiC are favorable, and its radiation stability needs to be validated. This implies that it is unlikely for thermal creep to limit the high temperature or lifetime of NITE SiC in fusion blankets. This still needs to be confirmed via thermal creep testing on NITE SiC samples. Its stability under neutron radiation is under special consideration due the potential in fission and fusion systems. The most influential criteria for

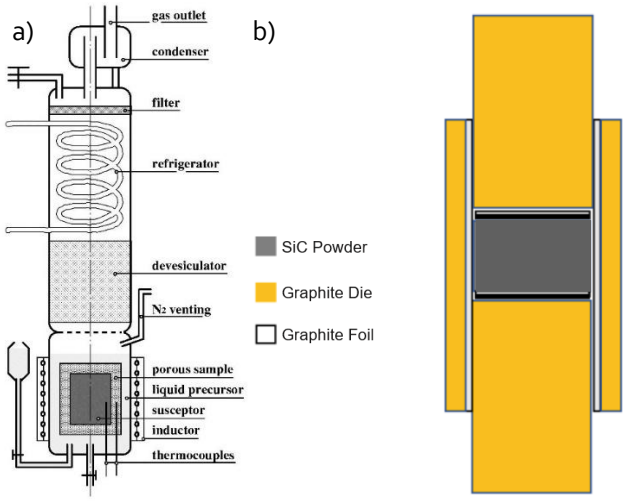

FIGURE 5. A) BOILING FILM CVI REACTOR WITH INDUCTION HEATING OF INTERNAL GRAPHITE SUSCEPTOR IMMERSED IN THE LIQUID PRECURSOR [10]; B) NITE SIC POWDER IN GRAPHITE DIE PREPARED FOR PRODUCTION.

radiation stability is the crystallinity of the material, which makes the ideal process seem to be CVD. Since CVD characterizations have been thoroughly examined, they were used as a benchmark to evaluate the performance of NITE SiC under neutron irradiation. It has been reported that up to 10 dpa irradiated NITE SiC had double the strength and 20% higher thermal resistivity when compared to CVD SiC [16]. The NITE route for producing SiC has been validated thus far, but if it is to be used in fusion environments, understanding how helium cavities form in the material is also necessary alongside neutron irradiation. **Figure 6** below shows the resulting swelling and defect resistivity recorded for both CVD and NITE SiC under neutron irradiation.

Fusion reactor materials must be able to remain functional despite neutron damage, as well as hydrogen and helium that will be produced by nuclear reactions within the material itself. The effects of He have been extensively studied in metals and alloys, while characterizations of many ceramics are widely unavailable. He implantation has been done for CVD SiC and other types of SiC composites, but not for NITE SiC. In a past study, He bubble showed preferential

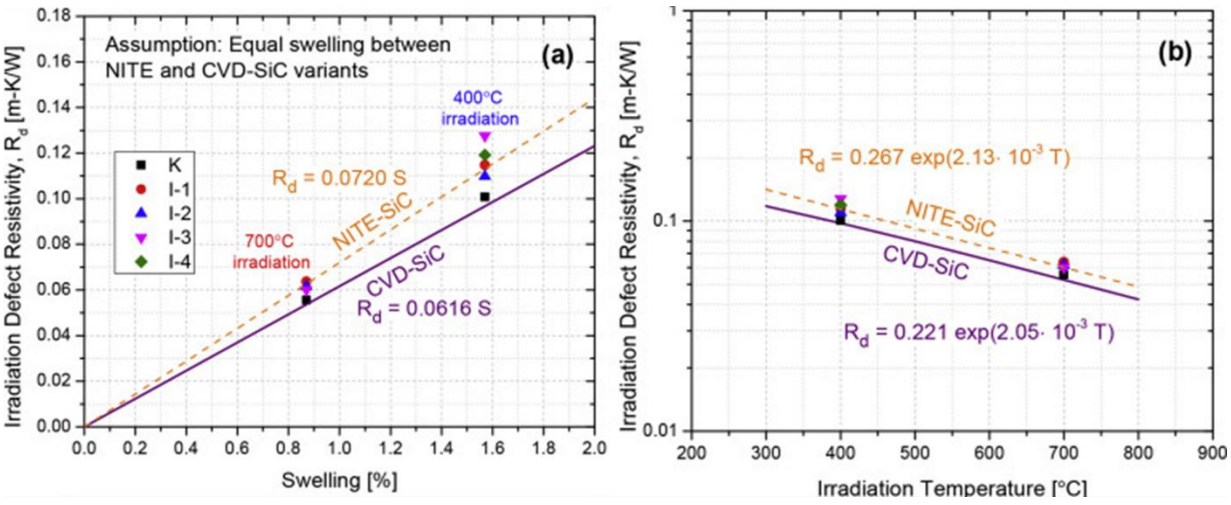

FIGURE 6. COMPARISON OF NITE AND CVD SIC SWELLING AND DEFECT RESISTIVITY UNDER IRRADIATION [16].

growth into the Si at Si/SiC boundaries [17]. This was done in a SiC/C composite which found no He along the graphite grains. Another study looked at 4H-SiC and found that at 1×10^{17} cm^{-2}, a decrease of about 50% in the hardness value is observed [18]. He migration during annealing also helps to reveal how much of the He the material holds onto as well as how much self healing it is able to go through at higher temperatures. It was found that at elevated temperatures, swelling increases with increasing fluence due to the damage accumulation and decreases with increasing implantation temperature because of the enhanced recombination of Frenkel pairs [19]. This study found that during annealing, the He bubbles concentrate along grain boundaries and form three-dimensional clusters of bubbles. This phenomenon become especially important when discussing NITE SiC as the additives (yttria and alumina) also accumulate in the grain boundaries around the SiC as a part of the production process. High-angle annular dark-field (HAADF) imaging has been used on NITE SiC. This kind of imaging shows heavier materials as lighter shades of grey so the rare earth metal oxides have been shown to create large gaps between SiC grains due to build-ups at triple junctions [16]. This study aims to contribute to the characterization of NITE SiC and how the He might interact with the oxide junctions. These characterizations heavily rely on microscopy techniques such as STEM, which includes HAADF, and TEM to accurately image structures and allow researchers to pull quantitative data from nanoscale imaging. **Figure 7** shows that in NITE SiC, fabrication additives tend to accumulate at grain boundaries (7a), which are also common migration sites for helium bubbles (7b), potentially leading to different bubble interactions than in additive-free SiC.

TEM operates by directing a high-energy electron beam through a specimen, typically less than 100 nanometers thick. As electrons pass through the sample, they interact with its internal structures, leading to scattering phenomena that carry critical information about material

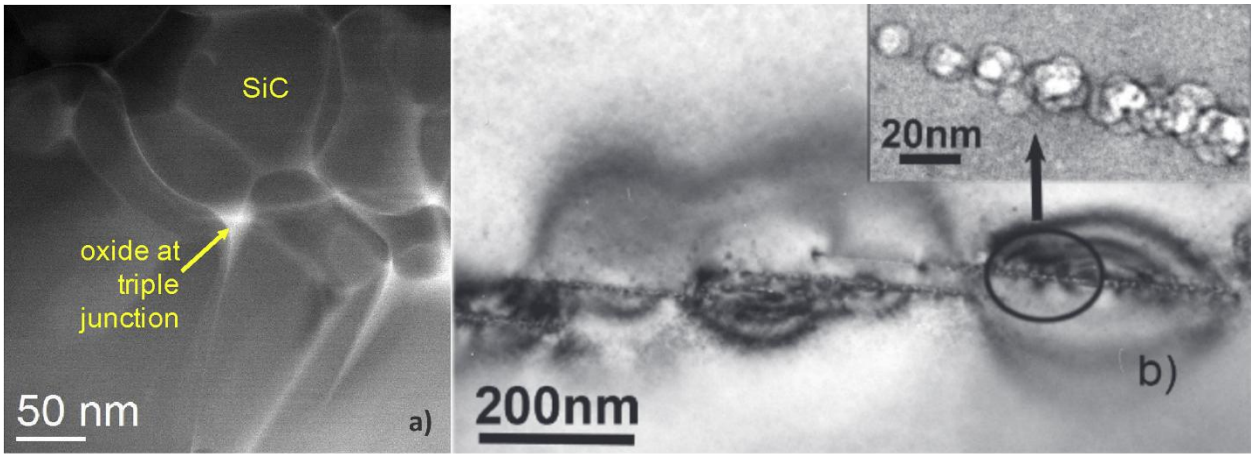

FIGURE 7. A) TRIPLE OXIDE JUNCTION IN NITE SIC [16]; B) HE BUBBLES COLLECTING ALONG SIC GRAIN BOUNDARY [19].

composition and arrangement. These transmitted electrons are then focused by objective lenses to form a magnified image on a viewing screen or detector. This process enables researchers to observe fine details such as crystal lattices, defects, and interfaces within advanced materials, providing invaluable insights into their properties and behaviors [20]. Since the data is received as a bunch of pixels, a lot of processing goes into microscopy so the data is in a form that can actually be utilized such as extracting the size or density of the desired features, in this case bubbles. The processing is currently set up so that one person has to go through all of the images taken and prioritize which pictures have the information that is necessary for characterization. A 1 hour experiment can result in 16,000 plus images. The experiment with 16,000 images had 50 of those images annotated so a lot of data is currently lost just due to the sheer quantity of images. Once the data is being annotated, one person has to use as software such as ImageJ to go through and individually mark each bubble. One image that was done for this study had 1,260 bubbles and took over six and a half hours to annotate by one undergraduate. The data then had to be fed through a custom Python file to calculate and extract the size of the bubbles observed. For an entire data set, this process can take weeks even when multiple people are working on it.

This leads to some data sitting for years before it can be analyzed. What good is data that is just sitting around? This is why machine learning is so interesting to the materials field. ML models such as the Segments Anything Model (SAM) created by Meta, are already capable of identifying anything from cars to cake to people. If applied to structures that are meaningful to the materials science field, there is potential to greatly shorten the amount of time it takes to analyze microscopy data and hopefully spare future undergraduates the pain of manually annotating data over the course of hours. Even if the model is only able to get 1/3 of the data from an image, if applied to the entire data set this might be able to counter balance any missed features. This is also really important to keep in mind in terms of systematic error. If it is not seeing the smallest of the bubbles we might introduce bias into the data set. This needs to be kept in mind when determining which ML model to use as well as what tools one uses to run said models. **Figure 8** below demonstrates the simplified principles of TEM and STEM microscopy and how the imaging techniques work.

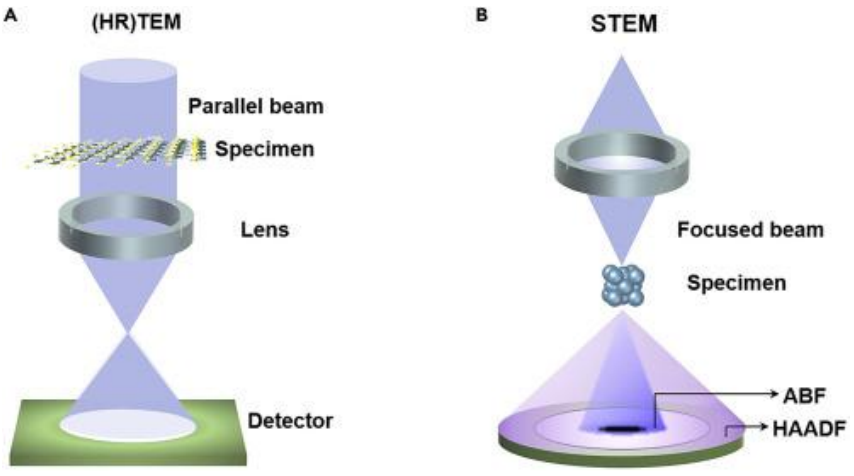

FIGURE 8. SIMPLIFIED SCHEMATIC OF A) TEM AND B) STEM SAMPLE SETUPS [20]

Data Analysis via Machine Learning

The field of microscopy is ever advancing and has come a long way from producing a few micrographs with a reported "calculation by Swann around 1970 estimating that all TEMs, since they first became available commercially (15 years), had only examined 0.3 mm^3 of material!" [21]. Nowadays, TEMs can produce very large amounts of data on the order of terabytes per day. These can result in thousands of frames and experiments that were once limited by data are now limited by the experimenter's ability to analyze and make use of the data [22]. There are many approaches to ML with distinct advantages and drawbacks to each method. K-means clustering is one such strategy that is purely statistical in nature; thus, does not require prior information about the material and it is widely available, which could lead to wider collaboration and development of materials science models. The disadvantage to this kind of model is that it takes extensive human knowledge to check model validity and each new material requires the same level of investment as the model has to be retrained each time [23]. While a statistical approach has not yielded the desired results, deep learning has been obtaining more success in TEM video processing, such as image quality improvement through convolutional neural networks (CNN) [24]. For this reason, a neural network based approach was selected for this experimental work. **Figure 9** below demonstrates the use of a trained CNN and the accuracy of the model predictions compared to a human annotator.

One of the major computational tools employed in this study was the YOLO (You Only Look Once) family of object detection algorithms. YOLO is especially well-suited for analyzing TEM micrographs due to its unique combination of speed, accuracy, and real-time inference capability. Unlike traditional object detection models that require a two-step process—region proposal followed by classification, YOLO performs detection in a single neural network pass,

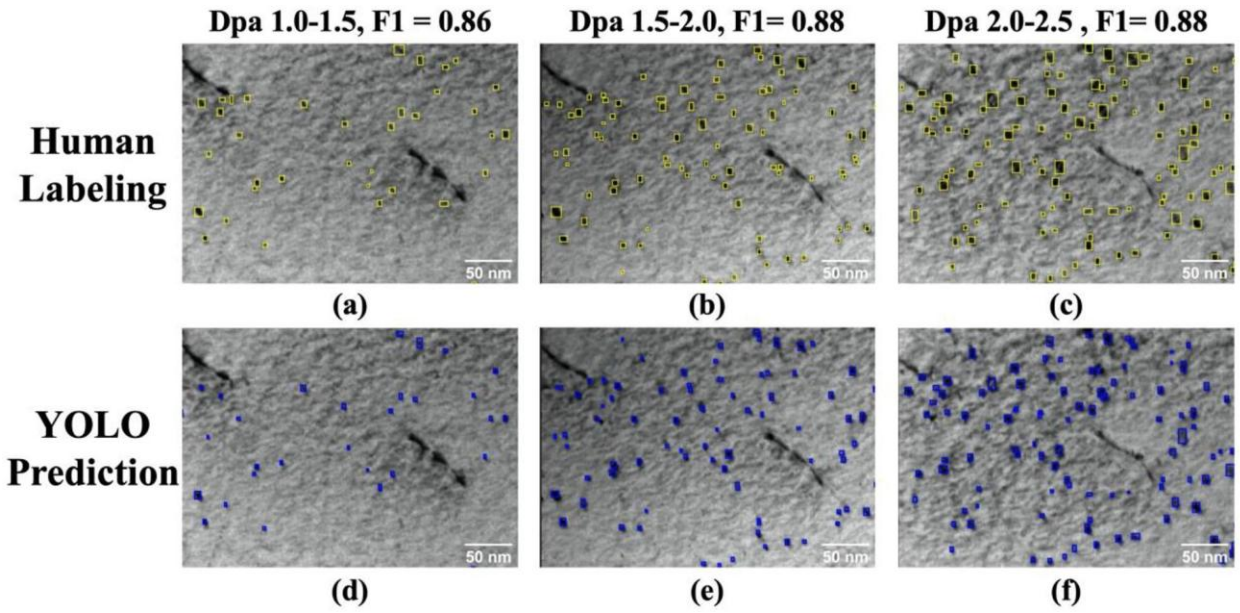

FIGURE 9. HUMAN LABELING COMPARED TO A YOLO BASED PREDICTION [22].

significantly accelerating image processing. This speed is crucial for large materials science datasets, particularly in high-throughput or *in situ* environments where real-time analysis may one day be required. In this study, the YOLOv8-S architecture was selected for its efficient balance between performance and computational demands. YOLOv8 represents a major evolution in the YOLO family, integrating advancements in backbone structure, training flexibility, and object localization precision, making it highly compatible with complex and noisy image data such as helium bubble micrographs [25]. **Figure 10** below shows the difference between single and two stage detectors to specify the differences between YOLO and SAM.

The application of YOLO to materials science has already shown strong precedent. Previous studies have demonstrated the utility of YOLO in automated defect detection during *in situ* ion irradiation experiments, reinforcing the model's relevance for real-world microstructural analysis [22]. The ability to track frame-by-frame features positions it as a compelling candidate for future dynamic microscopy studies, especially in fusion-relevant research where defect

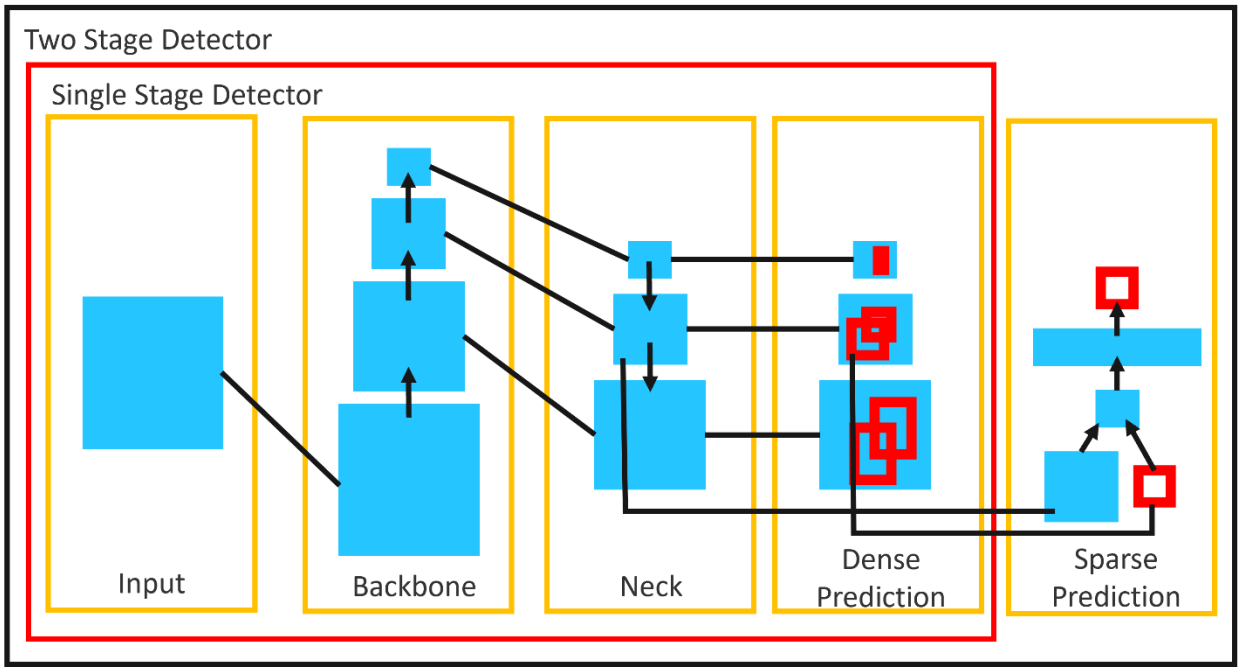

FIGURE 10. OBJECT DETECTOR ANATOMY FOR ONE AND TWO STEP IDENTIFICATION TECHNIQUES [25].

evolution must be monitored under changing environmental conditions. Despite these strengths, however, neural networks like YOLO remain notoriously difficult to interpret. Unlike human-defined algorithms, neural networks learn "rules" internally that are not explicitly programmed, and these rules are often highly abstract, making post-hoc reasoning challenging. This issue has been recognized since the early 1990s, and a past study highlighted the difficulty in extracting human-understandable logic from trained networks [26]. More recent efforts, use heatmap visualizations and relevance propagation techniques to better understand what neural networks are actually learning during training. More recent efforts aim to bridge this gap employ visualization techniques such as sensitivity analysis, deconvolution, and Layer-wise Relevance Propagation (LRP). Sensitivity analysis involves systematically perturbing input values, such as pixel intensities in an image, to see how much they affect the network's output, thereby identifying which parts of the input the model is most "sensitive" to. Deconvolutional methods

reverse the forward pass of the neural network to map activations back onto the input space, producing heatmaps that suggest which image regions contributed most strongly to the final prediction. LRP, on the other hand, distributes the network output backward through the layers to assign a relevance score to each input feature, highlighting areas the model "considers important" for the decision made [27]. These techniques reveal that models sometimes rely on seemingly irrelevant or unintuitive features to make decisions, an issue that is particularly problematic in scientific contexts where interpretability is as important as accuracy. **Figure 11** below shows the heat maps that resulted from the visualization techniques described above.

Given these limitations, the evaluation metrics for this study were intentionally chosen to include both statistical performance and time efficiency. The statistical component allows us to quantify the ability to generalize, detect bubbles of varying sizes, and perform across different material types with the models. The time efficiency evaluation speaks directly to the practical impact of automation in the research workflow. For example, annotating the dataset for the largest model in this study, Model 20, required more than 60 hours of manual effort. Reducing

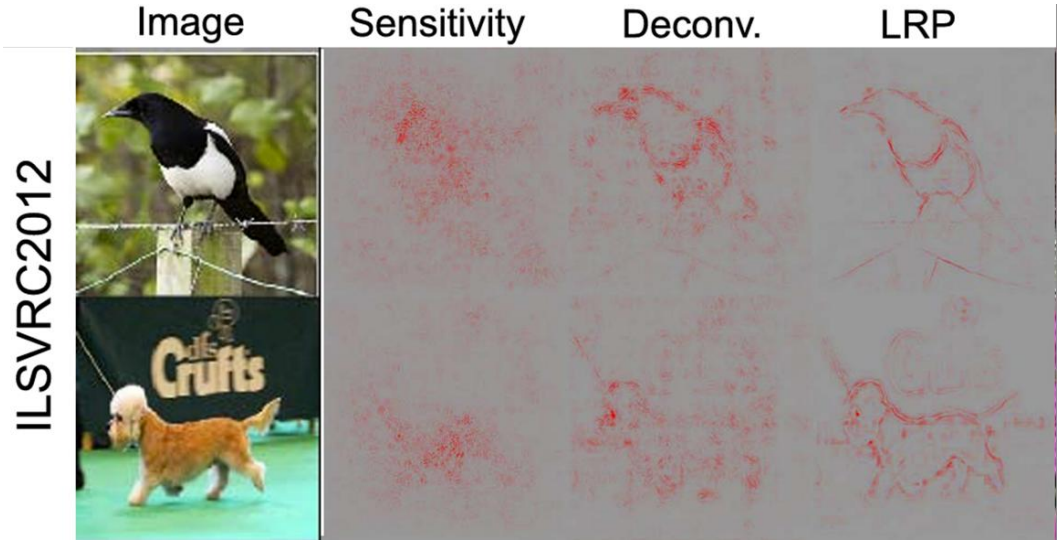

FIGURE 11. GNN EXAMPLES OF HEATMAPS FOR IDENTIFICATION OF ANIMALS IN IMAGES [27].

this burden through a high-performing model not only accelerates research timelines but also reduces cognitive fatigue and user bias, making materials characterization more consistent and scalable. Together, these evaluation strategies provide a balanced view of model utility, not just as a technical achievement, but as a transformative tool in modern materials science.

Chapter 3: Methodology

Model Development

Manual annotation analysis of the raw TEM micrographs of the irradiation-induced helium bubbles was conducted with ImageJ [28] and VGG Annotator software [29]. ImageJ software was used to determine the pixel to nm ratio. VGG Annotator was used to create manual annotations and saved in a JavaScript Object Notation (JSON) files. The annotations were then used to determine average size and eccentricity of the bubbles using a custom-built Python script. The script utilized the coordinates denoted by VGG to calculate the area and eccentricity for each annotation. Once these annotations are collected, they can be fed to the model to begin training so that automated annotation can be performed. The machine learning analysis of the same TEM micrographs, as well as a that of a single crystalline sample, were performed with the real-time microscopy image analysis Theiascope^{TM} platform [30] which ran Segment-Anything Model (SAM) and YOLOv8 machine learning models. The main aspect of the machine learning models investigated in this thesis was the effectiveness of the ML models when compared to manual annotation in terms of time. The SAM model is a "zero-shot" model that was trained on millions of images to segmentation every object in an image. It does this by looking for texture changes or color differences in pictures. SAM is able to identify helium bubbles, but with a lot lower confidence than a domain specific model for TEM-based material characterization for helium bubbles in SiC. SAM is freely available and requires no training, thus it was used to benchmark domain specific AI models developed using the manually annotated images. **Figure 12** below demonstrates the results of a manual annotation of 1260 bubbles in a lithium aluminate sample micrograph.

Ultralytics is a software company that develops and maintains variants of the You Only

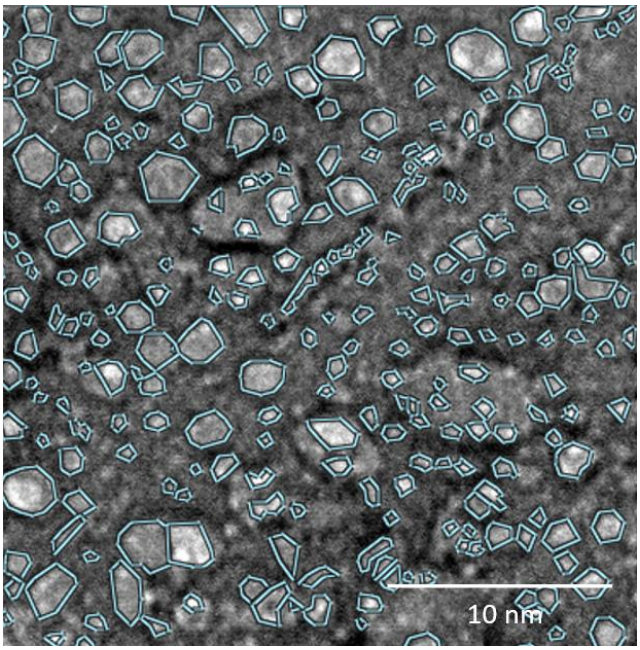

FIGURE 12. RESULTS OF MANUAL ANNOTATION OF MICROGRAPH RESULTS.

Look Once (YOLO) family of models. They provide an open source package, Ultralytics, for training YOLO AI models. This type of model is valued for its low latency and integrated tracking abilities, which is ideal for real-time object tracking in videos and potentially *in situ* TEM experiments. YOLOv8 was specifically chosen because it is included in the Ultralytics Python package, can performance pixel-wise instance segmentation, and includes two integrated tracking algorithms. A total of twenty-six YOLOv8 models were trained for this analysis. The variables changed for this included the size of the model, number of epochs, amount of input data, and whether or not image augmentation was used. The sizes included by Ultralytics are Nano, Small, Medium, Large, and X-large. The size of the model can be conceptualized as a mammalian brain which has more neurons as the size increases. The smaller models are ideal for resource-constrained deployment environments, such as mobile or embedded devices, like a "smart" camera. This in theory means that the X-large should perform better than the Small, thus this was investigated to see how the accuracy improved while taking the longer training time that

running larger models entails into account. Larger models also put more strain on the hardware being used to train it so there are some hardware limitations on the larger model sizes. Epochs can be thought of as one training cycle. As the number of epochs increases, the model has more and more chances to be more accurate. This can be conceptualized as playing darts when one is out of practice. At the beginning, the accuracy of the throws seems almost random. As more throws, epochs, are conducted the model can hone the process and become more accurate. Just as with darts, there is a diminishing return as the points, accuracy, will plateau at some point making more practice an inefficient use of time. Epochs were investigated to quantify how many epochs produce a near human result so that the time it took to train could be considered. The amount of input data refers to the total number of manual annotations fed to the ML model. This was investigated to determine a minimum amount of annotations for the accuracy desired for the experiment to ascertain how much time needed to be invested up front. Image augmentation is used to create a large number of images out of a small number of images for ML model training. One model was made with un-augmented data so that it could be determined how it compared to the augmented data and justify the extra time it takes to augment images. Model 5, 10, and 20 refer to the number of images fed to the model, but the number of bubbles annotated have also been listed for clarity. The time to annotate and time to train the models are also provided to contribute to the time efficiency comparison. All 26 models and their configurations are detailed below in **Table 1**.

The Ultralytics training process returns results for both precision and recall. Precision being a check that the bubbles identified by the model are truly bubbles (false positive) and recall being a check of the bubbles correctly identified what the total percent of bubbles detected is (false negative). The included models show pronounced trends with high precision and the best

TABLE 1. CONFIGURATION OF ALL MODELS TRAINED AND STUDIED.

Variables	Model 5	Model 10				Model 20
Bubbles	6779	9666				12661
Time to Annotate (hr)	33.89	48.33				63.305
Epochs	100	10	100	1000	100	100
Model Size	All	All	All	All	Medium	All
Augments	Yes	Yes	Yes	Yes	No	Yes
Time to Train (hr)	0.2232	0.0246	0.3494	3.6038	0.034	0.58833

recall of the different epoch models investigated. The precision and recall metric for a human annotator is considered 0.8. The results indicate that once you get to about 320 epochs, the difference between the different size models is negligible. The Small size model even pulled ahead of the other sizes which contributes to the argument that the size of the model does not matter for all data sets as they all plateau around 0.75. The recall is not at a satisfactory level, but this can be explained by the training process. The images that were annotated had an image size of 790 x 790 pixels and were compressed to 640 x 640 image size for training. This was done because YOLO is sensitive to image size and 640 x 640 is the default size used. The size was left at the default setting so that the base model could be evaluated. This compression makes the smallest of the bubbles even less detectable by the model. This indicates a need to investigate size and precision and determine a cut off of what is the smallest annotation, in pixels, the Ultralytics training process can see for future training procedures. These training results were used as a checkpoint to determine quickly if the training had been a success or not. If the

precision value was below 0.5, the training was considered a failure, and the inputs were evaluated to determine what caused the poor training. **Figure 13** below demonstrates the precision and recall of all of the 1000 epoch models; the low recall is indicative of a segmentation training error that would be corrected after this work and can overall be ignored.

If the models passed the check described above, then they would move onto the testing phase. Six test images were selected. Five of these images were from similar data sets that also demonstrated white He bubble cavities on a black background. The final image in the test set was one used in the training to determine the effectiveness of the model on its intended database. The models evaluated the image using Theiascope™ software from Theia Scientific. Count, confidence interval, and size of the segmentation masks produced by the models were collected

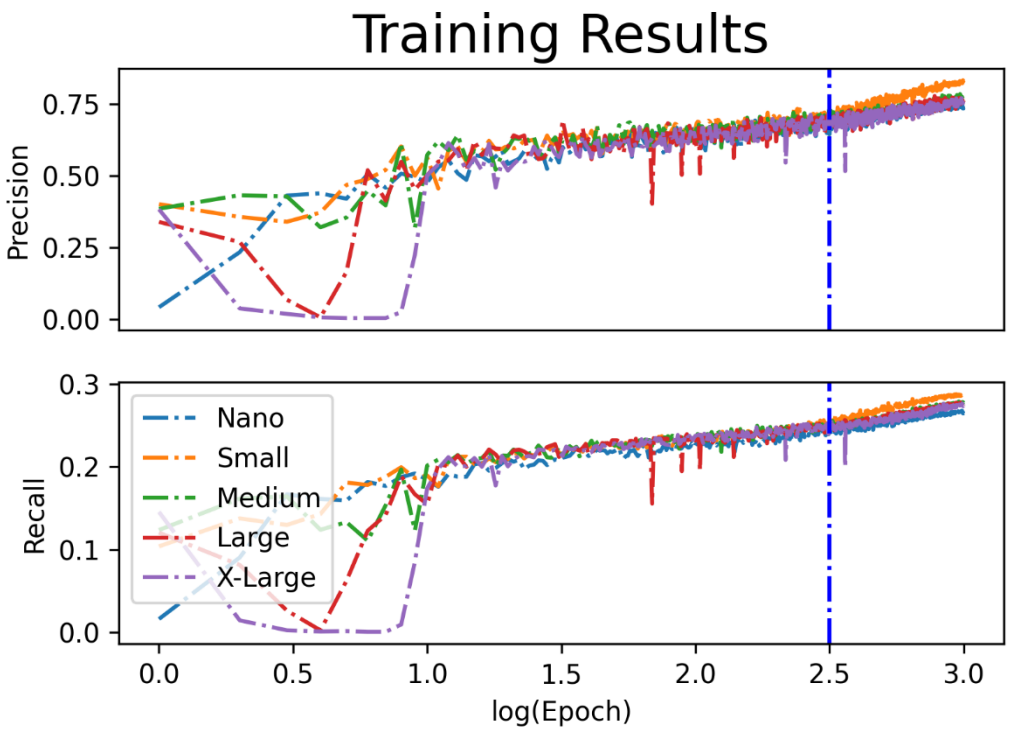

FIGURE 13. TRAINING RESULTS OF MODEL 10 WHEN EPOCH WAS 1000; THE BLUE LINE REPRESENTS EPOCH 320.

for evaluation of how the different variables affected the outcomes as well as how much time it took to accomplish.

The performance of the YOLO-based models, particularly their behavior on helium bubble detection across varied images, informed how the experimental process would be structured and evaluated. Understanding the limitations of the model, especially the challenges the model faced with smaller features and unfamiliar data types, highlighted the need for a controlled, comparative experiment using well-defined materials. This directly influenced the decision to adapt a previously established helium implantation methodology and apply it not only to conventional CVD SiC but also to NITE SiC, which had not been included in the original study [31]. Including both materials allowed for an intentional contrast in microstructure and mechanical behavior, offering a more meaningful dataset for evaluating model generalization and segmentation consistency. In this way, the model development efforts and computational insights fed directly into the experimental design, ensuring that the materials chosen and the analysis performed could serve as both a validation and a stress test of the proposed automation pipeline.

Experimental Methods

NITE SiC was fabricated by mixing a composition of three powders of silicon carbide (500 nm, AB147382, Lot 5101, HC Starck GmBH), aluminum oxide (80 nm, US3008, US Research Nanomaterials), yttria oxide (45 nm, US3551, Lot 1314-36-9, US Research Nanomaterials) in a mass fraction of 94:3.9:2.1, respectively, and PEI at 1 wt% of the total powder mass. The mixture was blended at 1600 RPM for 3 minutes (Speedmixer, FlackTek, Inc, Landrum, SC, USA) into an ethanol slurry at a 33% solid loading before being dried. X-ray diffraction (Aeris XRD Diffractometer, Malvern Panalytical, Ltd, Malvern, UK) was conducted

and confirm the recorded mix ratios are likely accurate, validating the previously recorded phases and the XRD results can be seen below in **Figure 14**.

The samples themselves were made via spark plasma sintering with pressure assistance these were fabricated with the same method used in previous SiC component experimentation done at the High Temperature Materials Laboratory at the University of Tennessee [32]. The three samples used for this experiment were labeled 147, 150, and 151 and were fired at 1825 ^{o}C and held there for 5 minutes at 10 MPa. These samples where made using cylindrical graphite dies and resulted in a circular wafer shape about 13 mm in diameter made using previously established techniques [33]. High purity SiC samples were machined from CVD material from Rohm and Haas Advanced Materials (now Dow Chemical Co., Midland, MI, USA). The samples were cut into fourths with a Pace Technologies PICO155P precision cutter diamond saw. The NITE SiC samples were cut at 1500 rpm while the CVD SiC samples were cut at 1200 rpm. The CVD SiC samples had been previously polished and were kept as is. The NITE SiC samples

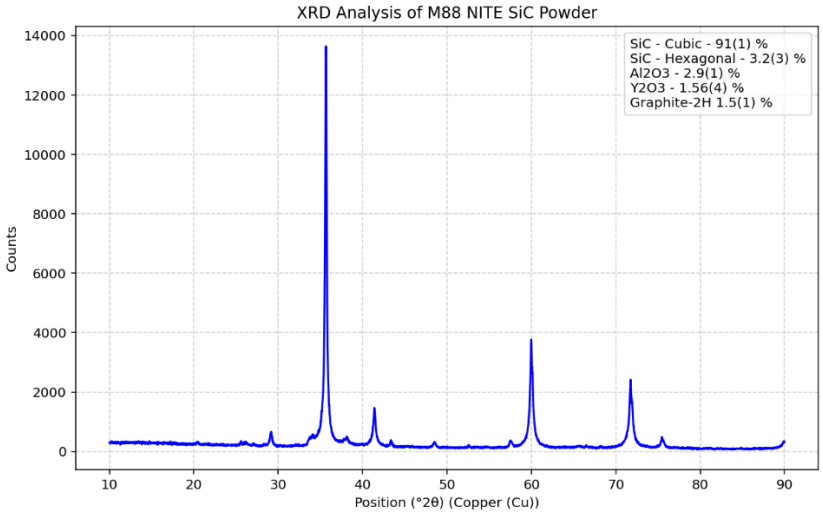

FIGURE 14. X-RAY DIFFRACTION ANALYSIS OF MIX 88 OF NITE SIC.

were mounted in epoxy and polished to 1 µm with a Pace Technologies NANO 1000T grinder-polisher. **Table 2** below lists the fabrication parameters for the NITE SiC samples and their resulting densities which are around 95% fully dense (theoretical density is 3.21 g/cm^3).

The helium implantation was performed in conjunction with the Ion Beam Materials Laboratory (IBML) through the Center for Integrated NanoTechnologies (CINT) user proposal. The implantation procedure of a similar study done on CVD SiC which investigated how He cavities form in the material at two different fluences was used as a guideline for this experiment [31]. This study involved ion irradiation and a He implantation at 277 ^{o}C with 65 keV He ions. The implantation was performed with the implanter at the IBML at Los Alamos National Laboratory (LANL) with the same conditions to three different doses instead of two. This was done to fluences of 1 x 10^{14}, 1 x 10^{15}, and 1 x 10^{16} ions/cm^2. Each dose would be performed on 2 CVD and 3 NITE SiC samples for a total of 15 samples at 3 unique doses. Due to the large sample holders found at this facility, all the samples for each respective dose were done at once.

TABLE 2. SAMPLE FABRICATION SPECIFICATIONS FOR SPS METHODS.

Sample	Firing Temperature (^{o}C)	Firing Time (min)	Mass (g)	Bulk Density (g/cm^3)	Accupyc Density (g/cm^3)
147	1825	5	2.063	3.187	3.154
150	1825	5	2.001	3.188	3.169
151	1825	5	2.004	3.195	3.165

Figure 15 below demonstrates the irradiation conditions with approximate dimensions of the samples and provides images of what the CVD and NITE SiC samples look like pre-irradiation.

The samples were used for mechanical evaluation. The mechanical stability was evaluated via nanoindentation (NMT04 NanoMechanical Testing System, FemtoTools AG, Buchs, Switzerland). A 40 μN force with a Berkovich tip was used with a maximum depth 0.25 μm in force controlled mode using CSM indentation. The amplitude and frequency associated with the indents were 0.002 μm and 45 Hz. A 10 by 10 array was done on all 15 samples for a total of 100 measurements each for statistical accuracy.

The mechanical testing process was designed not only to capture reliable baseline hardness values but also to establish how each material responded to controlled helium implantation. Introducing helium into the SiC samples simulated the radiation environment expected in fusion applications, where helium is produced as a byproduct of fusion reactions and

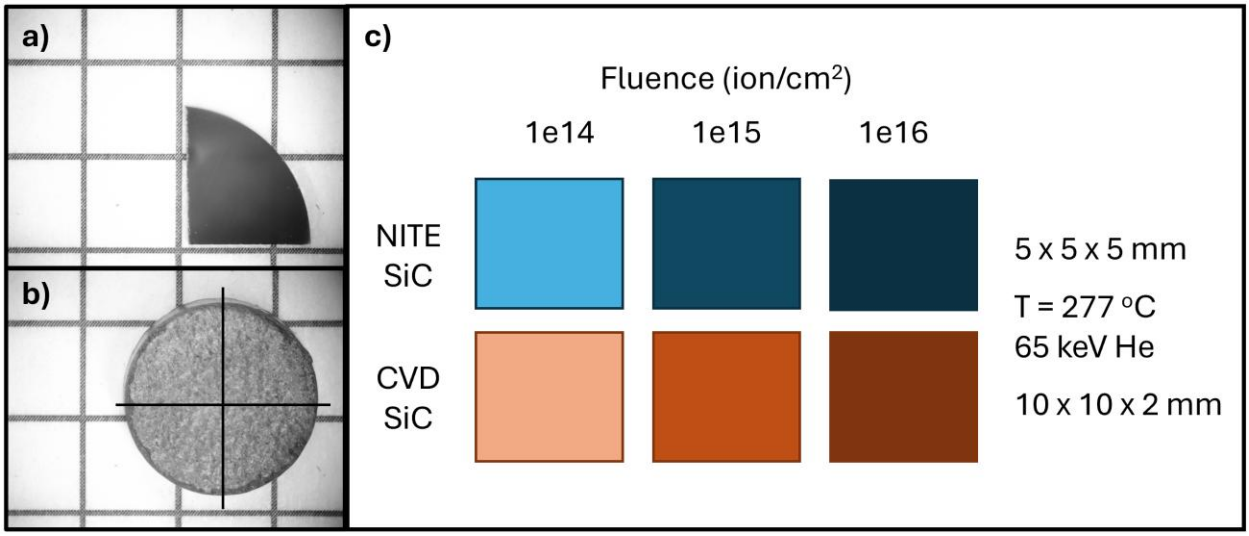

FIGURE 15. SAMPLE SET A) EXAMPLE OF CVD SAMPLE WHICH HAS BEEN QUARTERED AND POLISHED (THE SQUARES ARE 0.25" OF 0.635 CM); B) EXAMPLE OF A NITE SAMPLE BEFORE BEING CUT AND POLISHED; C) WITH CORRESPONDING MATERIALS, DOSES, AND IMPLANTER CONDITIONS.

contributes to long-term material degradation. Three distinct helium implantation doses were used to observe how different exposure levels might influence mechanical behavior. This approach enabled a systematic investigation into dose-dependent effects on hardness and potential embrittlement. By applying a consistent array of nanoindentation measurements across all samples, the experimental design ensured that any differences observed could be confidently attributed to either the material type or the implantation conditions. These results would later form the backbone of comparisons with microstructural imaging and model-based bubble analysis, creating a complete picture of how helium affects SiC across both physical and computational dimensions. **Figure 16** shows the IBML in which the irradiation was performed.

A simulation was done via the stopping and range of ions in matter (SRIM) with the same specifications as the implantation to gain an estimated depth that the He ions would reach. The density of the SiC target was assumed to be the theoretical density of 3.21 g/cm^2 [5] and the resulting range was between 200 and 400 nm with a concentration at 320 nm. This data will later be used when samples are prepared for microscopy and can be seen below in **Figure 17**.

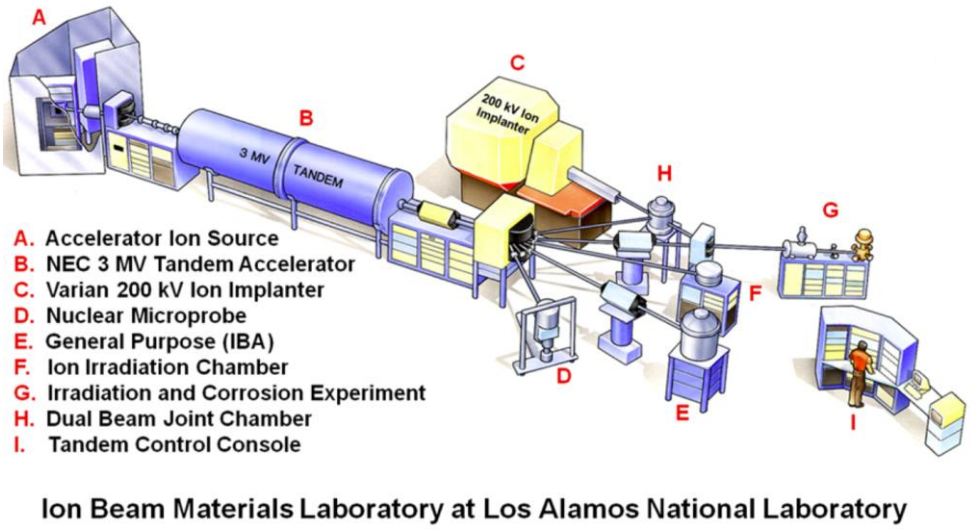

FIGURE 16. IBML AT LANL USED FOR HE IMPLANTATION [34].

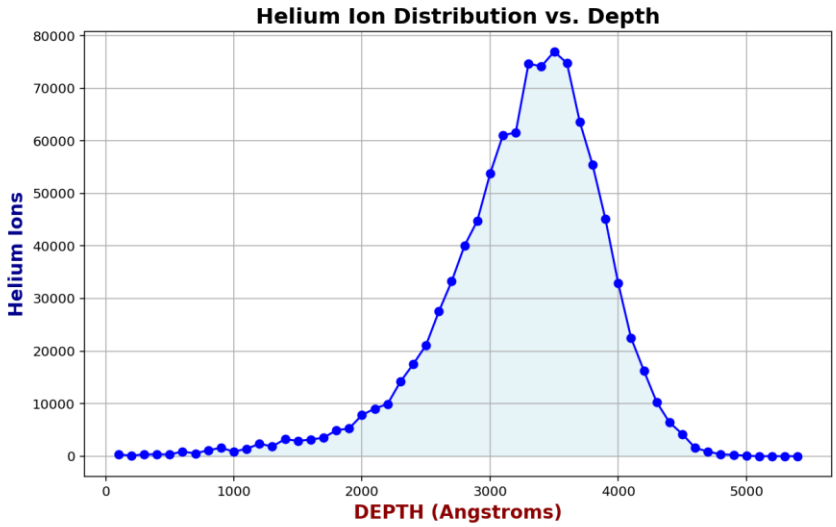

FIGURE 17. SRIM SIMULATION OF HE IONS AT 65 KEV IN SIC AT THE THEORETICAL DENSITY OF 3.21 G/CM^3 SUGGESTS THAT 319 µM IS THE OPTIMAL DEPTH FOR MICROSCOPY [35].

Chapter 4: Results

The following section begins with an overview of the preliminary model results, focusing on their performance metrics and what those outcomes suggest about the reliability of the model and applicability to helium bubble detection. This analysis provides a foundation for evaluating the broader potential of machine learning tools in materials characterization workflows. The following segment then transitions into the experimental findings, specifically the nanoindentation data collected from both NITE and CVD SiC samples. Together, these sections aim to connect computational insights with physical measurements, highlighting how each contributes to a deeper understanding of material behavior under helium implantation.

Model Statistics

To evaluate the efficiency and accuracy of various segmentation approaches, a comparative analysis was conducted between manual annotations, the Segment Anything Model (SAM), and multiple YOLOv8 configurations. The manual annotation consisted of 1260 bubbles that took approximately 6.45 hours to complete. The segmentation mask results from the manual annotation, SAM annotation, and two YOLOv8 models that had the highest counts are shown below. The SAM model took no training and was able to identify nearly as many bubbles as the Medium YOLO model with 1ooo epochs that took 3.579 hours to train. SAM did not correctly identify many of its bubbles; however, the Figure 18b below demonstrates that there were a lot of overlapping bubbles that were not counted and separate bubbles in the manual annotation. This indicates that SAM is less efficient than initially thought when used directly on the micrographs. Despite this, the only models that had higher counts than SAM were the Small and Medium YOLO models with 1000 epochs. This indicates that the higher end of our counts did not go high enough to find the plateau in the time to benefit ratio. It also demonstrates that the

size of the model does not affect the performance of the model as much as the number of epochs do. Taking this information into account, further statistical investigation will focus on larger numbers of epochs and stick with the Small sized YOLO model. A comparison of the performance of the manual, SAM, YOLO Small, and YOLO Medium annotations can be seen below with their relative time or training time for YOLO in **Figure 18**.

The increasing epochs and data input investigation involved testing the individual models on 5 separate test images and having the counts, confidence score, and area averaged to be compared to the manual annotation. The results of the varied size and epoch models can be seen below. The models started with varied count results with the expected result of the X-Large model performing the best with the highest number of counts. As the epochs increase, the models all have increasingly similar results, and some begin to identify more bubbles than the manual annotation at 1000 epochs. The confidence score results follow much the same trend as the counts with all of the models settling around 80% confidence at 1000 epochs. The area results are about the same trend, but they never quite reach the average found via manual annotation. This indicated that all the models, no matter the size or epoch, struggles to find the smaller bubbles which contribute to a higher average area size and would insert a bias in an evaluation of real data using the models produced. This further supports that the size of the model matter much less to the success of the model than the number of epochs. The results can be seen below in **Figure 19**.

The increasing data input investigation was performed on the same 5 test images as the epoch tests were. The results for Model 5, Model 10, and Model 20 can be seen below. The models were fed 6779, 9666, and 12661 bubbles, which took 33.895, 48.33, and 63.305 hours respectively. Unsurprisingly, Model 20 which had the most data performed the best in all three

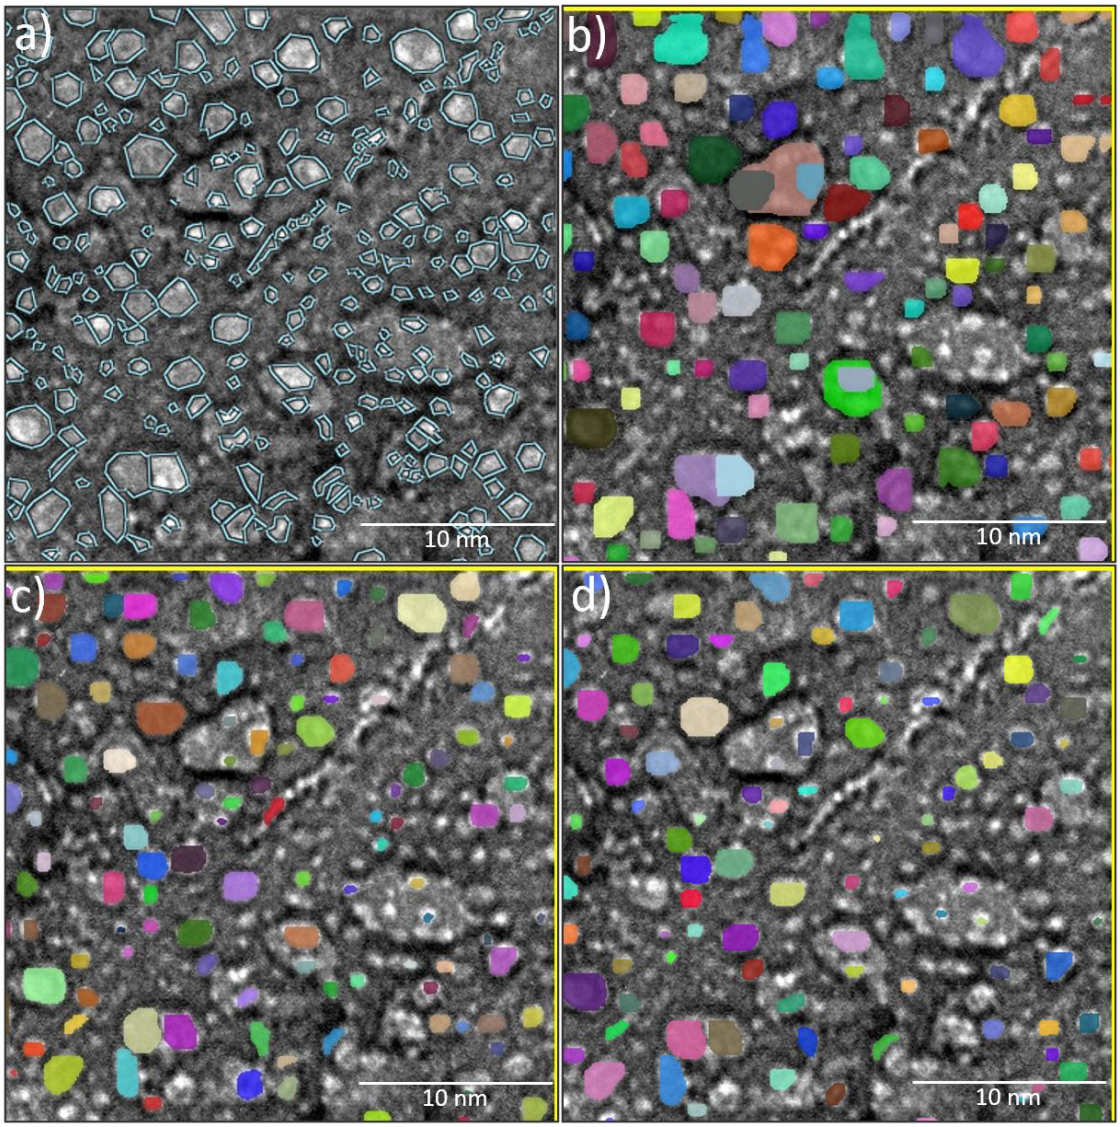

FIGURE 18. A) MANUAL ANNOTATION (1260 BUBBLES, 6.45 HOURS); B) SAM ANNOTATION (349 BUBBLES, 6 SECONDS); C) YOLOV8 SMALL, EPOCH = 1000 (403 BUBBLES, 3.555 HOURS); D) YOLOV8 MEDIUM, EPOCH = 1000 (354, 3.579).

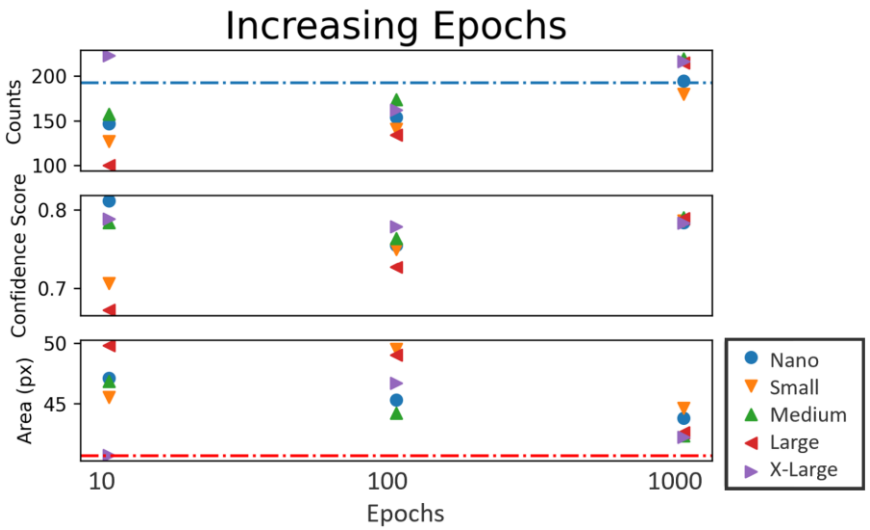

FIGURE 19. EVALUATION AS THE NUMBER OF EPOCHS INCREASE FOR MODEL 10. THE BLUE LINE IS COUNTS (192.6) AND THE RED LINE IS AREA IN PIXELS (40.67) WHEN MANUALLY ANNOTATED.

evaluated categories. Model 20 was the only model that was made with two different data sets, which had helium bubbles in two different materials. This suggests that new models do not have to be made every time a new experiment is run and that one large helium bubble detecting YOLO model could be made. Once the model reached a certain level of success it could be applied to future experiments with no need for more manual annotation. This prompted an expansion of the original goals for applying a model to the planned He implantation in SiC. The data from increasing the input annotations can be seen below in **Figure 20**.

The model analysis provided valuable insight into the trade-offs between training time, detection performance, and annotation efficiency across different architectures. These findings not only informed the selection of parameters for further automated analysis but also highlighted the practical limitations of current segmentation tools when applied to helium bubble detection in

TEM micrographs. While the computational results established a framework for scalable image analysis, understanding how helium implantation physically alters the material remains

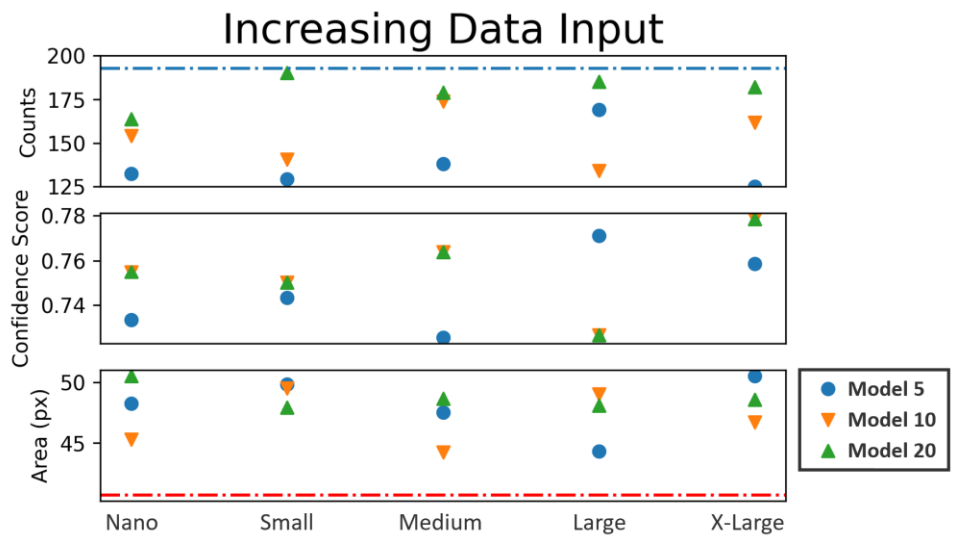

FIGURE 20. ACCURACY EVALUATION AS THE INPUT DATA INCREASES. THE BLUE LINE BEING THE AVERAGE COUNT WHEN MANUALLY ANNOTATED (192.6). THE RED LINE BEING THE AVERAGE SIZE WHEN MANUALLY ANNOTATED (40.67).

essential. The next section presents the mechanical evaluation of the helium-implanted SiC samples, offering a complementary perspective on material performance that will support and contextualize the image-based findings.

SiC Mechanical Testing

The nanoindentation data revealed that the CVD samples were consistently 15% harder than the NITE samples for every fluence tested. The CVD samples saw a 4% increase between 1 x 10^{14} and 1 x 10^{15} ions/cm^2. The NITE samples had relatively large error margins, which indicates that further testing should be done to confirm these results. The resulting hardness data can be seen below in **Figure 21**.

The mechanical results demonstrated slight trends in hardness values across the different helium implantation doses and material types, providing a quantitative foundation for understanding how each SiC variant responds to irradiation. These findings, when viewed alongside the segmentation and model performance data, offer a more comprehensive picture of both the physical and computational behaviors under study. With this context established, the

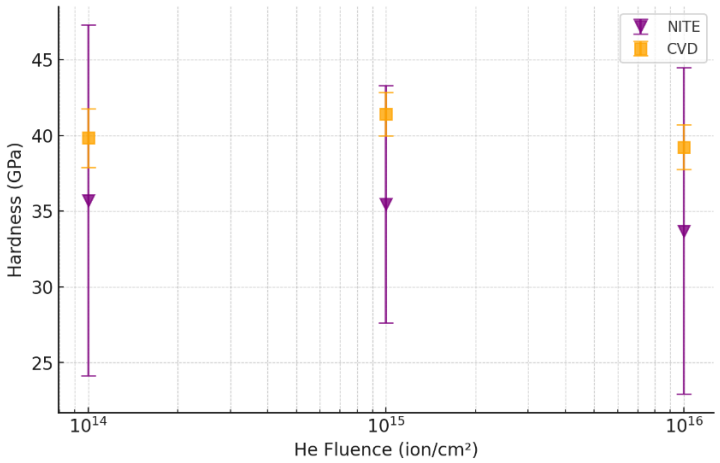

FIGURE 21. HARDNESS DATA COLLECTED FROM NITE AND CVD SIC SAMPLES.

following discussion will explore how these results align with expectations, where they diverge, and what implications they carry for future work in automated microstructural analysis and fusion-relevant materials research.

Chapter 5: Discussion

The findings of this study present several unexpected yet promising observations regarding the mechanical performance of CVD and NITE SiC samples. Contrary to previously published work, which generally showed that NITE-processed SiC composites performed better under irradiation in terms of toughness and crack resistance [16], the current unirradiated samples showed a higher average hardness in the CVD group. This divergence raises important questions about the intrinsic microstructural properties of each sample set, as well as the potential influence of fabrication parameters that vary across batches. The significant variability in the measurements, as evidenced by the large standard deviations, highlights the necessity of a larger dataset for a more statistically robust conclusion. Mechanical testing alone, without complementary imaging and microstructural validation, can offer only a partial picture of material performance. Moreover, the types of irradiation used in past studies were not always directly comparable. While earlier experiments may have involved neutron irradiation or different ion species, the current work focuses specifically on helium implantation. Helium, unlike other ions, has unique interactions with ceramic matrices, forming nanoscale cavities and pressurized bubbles that can weaken grain boundaries or trigger microcracks. This interaction is particularly important in fusion reactor environments, where He is produced *in situ* during deuterium-tritium fusion. Helium behavior in SiC is highly sensitive to variables such as temperature, crystallographic orientation, and local stress states, thus future work must incorporate a systematic, microscopy-based evaluation to elucidate the mechanisms responsible for bubble nucleation and growth. This deeper characterization has been delayed due to unforeseen equipment challenges, particularly in preparing TEM samples, but remains a top priority.

On the computational side of this study, the development and evaluation of the YOLOv8-based object detection model offered critical insight into how machine learning tools can be used to streamline and scale the analysis of helium bubbles. The initial statistical analysis of the training process revealed three key findings. First, while it might be expected that larger models (such as YOLOv8-X) would yield better performance, the results showed that model size had a comparatively small impact relative to the number of training epochs. Smaller models like YOLOv8-S were able to achieve nearly identical detection performance when trained over a sufficient number of iterations. This is a promising result for labs with limited computational resources, as smaller models are faster to train and require less hardware without compromising significantly on accuracy. The second observation was that the trained models consistently struggled with identifying small helium bubbles, especially those below a certain pixel threshold. This is a well-known limitation in computer vision systems, which are inherently biased toward detecting features that are spatially prominent (large). To address this, image tiling was implemented in subsequent experiments. Tiling breaks large images into smaller patches, effectively increasing the resolution of small features relative to the receptive field of the model. By forcing the model to operate on these smaller, higher-detail tiles, performance on similar data significantly improved. This strategy will likely remain an essential preprocessing step in future studies involving fine-feature recognition. The third insight from the modeling work was perhaps the most impactful: the trained model was capable of identifying helium bubbles in materials beyond the initial training set. This indicates that the features the model learned to recognize, contrast patterns, shape distributions, and spatial context, are generalizable across different microstructures. As a result, the focus of the project shifted from hyper-optimizing a specific model for one dataset to developing a broadly applicable “mega model” that could generalize

across experiments without the need for manual annotation. This opens the door to creating a foundation model for materials microscopy, similar in philosophy to the Segment Anything Model (SAM) from Meta that could be fine-tuned to new tasks with minimal data or training time.

The benefits of such a model extend well beyond research efficiency. In the course of this project, over 60 hours were spent manually annotating helium bubbles across a large dataset in preparation for model training. This not only introduces potential bias from human fatigue but also represents a major bottleneck in materials science workflows. By building a high-performing generalized model capable of inference on new datasets without user intervention, researchers can reduce the time from imaging to insight by orders of magnitude. Such improvements in analysis speed can accelerate the iterative process of hypothesis testing, material optimization, and reactor design. To support this ongoing development, a dedicated GPU computing system is planned for installation in the lab. This system will serve as a local resource for model training, experimentation, and deployment. While the final hardware specifications have yet to be finalized, it will enable much faster iteration on model architectures, hyperparameter tuning, and data augmentation pipelines. Additionally, it will allow real-time feedback for visiting researchers and facility users, fostering a collaborative environment where machine learning tools can be shared and co-developed. This system is foundational not only to the ongoing helium bubble analysis project but also to broader initiatives aiming to integrate AI into microstructural characterization across different materials and imaging techniques.

In parallel with the computational developments, this work also emphasizes the importance of revisiting the mechanical and microstructural properties of different SiC processing routes, particularly in light of fusion reactor demands. While the current hardness

data suggest superior strength in CVD-SiC samples, the picture is incomplete without understanding how these materials behave under stress, after irradiation, or at the interfaces where different materials must be joined. The fusion community has long recognized the difficulty of joining SiC components without introducing structural weaknesses or incompatibilities. Further investigation into SiC joint technologies, especially for NITE-based composites, remains a critical path forward due to real world applications requiring more complex shapes that cannot be created in one piece via sintering. Joining methods such as spark plasma sintering (SPS), active metal brazing, and transient eutectic bonding all show potential but come with trade-offs. For example, active brazing may provide strong bonds but risks contamination or embrittlement near the interface. SPS offers lower processing temperatures but may introduce thermal gradients that destabilize the microstructure. Transient eutectic methods aligned with the NITE process may offer a promising route due to better compatibility with the fiber matrix architecture, but still require more long-term durability studies under neutron irradiation. Understanding how helium bubbles and irradiation damage interact with joint interfaces is especially important. Helium can migrate preferentially toward grain boundaries or regions of stress concentration, which are often found at or near joints. If helium accumulation leads to blistering or decohesion at these sites, the structural integrity of reactor modules could be compromised. Incorporating joint analysis into TEM workflows and applying the same machine learning tools developed for bubble detection may help reveal failure mechanisms before they occur in service.

Altogether, this study offers a multi-layered approach to understanding and optimizing materials for next-generation nuclear energy systems. It combines classical mechanical testing with modern AI tools and outlines a roadmap for future research infrastructure, both in terms of

hardware and data pipelines. The work here not only challenges existing assumptions about SiC processing routes but also establishes a framework for scaling up data-driven materials science. By continuing to integrate detailed imaging, advanced modeling, and collaborative infrastructure, this research lays the groundwork for more rapid development of materials capable of surviving the extreme environments inside fusion reactors.

Chapter 6: Future Work

Due to technical difficulties associated with key equipment used in the sample preparation process for TEM, high-resolution microscopic imaging has not yet been completed on the CVD and NITE silicon carbide samples. However, the groundwork has already been laid, with the samples currently undergoing processing to ensure they are suitable for detailed microstructural analysis. This analysis will serve multiple purposes. First, it will allow for a direct comparison of helium bubble formation between the two fabrication techniques, enabling a quantitative assessment of bubble size, distribution, and density. Second, it will test the robustness and adaptability of the current machine learning-based micrograph analysis model by applying it to newly obtained data. The inclusion of this new dataset is expected to further train and refine the model, possibly revealing insights not present in the original image set.

This forthcoming round of TEM characterization will be critical in validating whether the trends identified in previously analyzed samples are consistent across different SiC processing techniques. If discrepancies are observed, this could point toward key differences in microstructural stability and irradiation resistance, which are essential parameters in selecting materials for fusion reactor environments. High-resolution imaging of both grain interiors and boundaries, as well as bubble nucleation sites, will help correlate processing methods with their effects on material performance.

Another major area of future work involves the creation of a larger, more generalized micrograph analysis model. Inspired by large-scale vision foundation models such as the Segment Anything Model (SAM) from Meta, this project envisions a machine learning tool that can be fine-tuned across multiple imaging modalities and materials systems. The goal is to build a model trained on a wide variety of microstructural images—across different materials, scales,

and experimental setups—so that it can learn generalizable patterns that are useful for detecting voids, inclusions, bubbles, grain boundaries, and other features across numerous domains in materials science. Such a model would eliminate the need to retrain networks from scratch for each new study, reducing development time and making advanced image analysis accessible to a broader community of scientists and engineers.

To support this vision, our facility is currently planning the installation of a local GPU-based computing system dedicated to the training and deployment of machine learning models for materials science applications. While the specific system has not yet been finalized, it is intended to provide high-throughput, on-site computational capacity to support both current and future AI research. This resource will enable faster training of large models and the real-time processing of high-volume microscopy datasets, which is especially important as new instruments produce increasingly detailed and complex data. Moreover, the system will be made accessible to external users of our facility, providing a valuable tool for collaborative research and user-driven model development.

In addition to modeling and hardware development, further investigation into the properties and performance of SiC joints is also a critical direction for future research. For SiC to serve as a reliable structural material in fusion reactors and other high-stress environments, effective joining methods must be developed and validated. Fusion reactors, in particular, require large monolithic or modular components to be bonded together without compromising thermal conductivity, radiation tolerance, or mechanical strength. Traditional joining methods, such as brazing and diffusion bonding, often fall short due to the inherent brittleness of ceramics and the mismatch in thermal expansion coefficients.

Emerging studies have explored several advanced joining techniques tailored for SiC-based composites. Spark plasma sintering (SPS) is another technique under investigation for joining SiC/SiC composites at relatively low temperatures and with minimal grain growth, preserving the original material properties [36]. Transient eutectic phase bonding, especially when aligned with NITE processing routes, may also offer potential in creating uniform and defect-free joints under fusion-relevant conditions [37].

Notably, NITE-processed SiC may present unique advantages for joining applications due to its inherent microstructure and densification behavior. The fine-scale, interconnected matrix in NITE composites offers improved diffusion pathways, which may enhance joint strength and reliability compared to CVD-SiC. Additionally, the mechanical flexibility and fiber reinforcement in NITE SiC/SiC systems could provide better resistance to thermal cycling and irradiation-induced stress, further emphasizing its value in next-generation reactor architectures.

To fully evaluate the performance of SiC joints, future studies must incorporate post-irradiation examination and high-resolution microscopy of joint interfaces. These efforts will be critical for identifying the formation of secondary phases, voids, or other structural irregularities that could compromise performance. Furthermore, integrating machine learning tools into the analysis of joint micrographs may streamline the identification of failure-prone regions, making joint characterization more efficient and predictive.

Ultimately, this multi-pronged future work—encompassing advanced TEM imaging, development of scalable machine learning models, deployment of dedicated computational hardware, and a deeper exploration of SiC joining techniques—aims to advance both fundamental materials science and applied nuclear engineering. These steps, which I plan to be

the focus of my doctoral research, will lay the groundwork for more robust, interpretable, and scalable solutions in the deployment of fusion energy and high-performance ceramic systems.

Chapter 7: Conclusion

A central concern in fusion development is the durability of materials exposed to extreme operational environments. Silicon carbide (SiC), particularly in fiber-reinforced composite form, has emerged as one of the most promising candidates for use in structural and blanket components due to its combination of thermal, mechanical, and radiation-resistant properties. Yet, the behavior of this material under high neutron flux, including its response to helium implantation and microstructural evolution, still demands deeper understanding. Electron microscopy offers invaluable insight into these behaviors, but the field is now confronted by a growing data bottleneck. With each experimental run generating thousands of images, traditional manual analysis methods are no longer sufficient. This has driven a shift toward automated, data-driven techniques that can process, analyze, and interpret materials characterization data at scale. The integration of machine learning, particularly object detection algorithms and neural networks, presents a significant opportunity to accelerate this workflow. These tools can not only identify helium bubbles and defects within SiC micrographs but can also uncover patterns that may elude even experienced researchers. The initial model performance and accuracy were much lower than any scientist would want to use for actual data. Despite this, the information gathered has been a boon for the development of materials sciened based machine learning. It also proved that one does not have to be a computer scientist to apply machine learning to other fields. Continuing to feed the model data and expanding its capabilities will reveal how useful these models can be at saving experimenters processing time.

In parallel, the advancement of SiC as a functional material in fusion reactors invites more detailed comparison between production methods. Chemical vapor deposition (CVD) and nano-infiltration transient eutectic (NITE) processes each offer distinct advantages and

limitations. While CVD yields highly pure, uniform material suited for demanding applications, NITE enables the cost-effective manufacture of fiber-reinforced composites with competitive performance metrics. Understanding the trade-offs between these two processing routes—particularly in terms of mechanical robustness, helium bubble resistance, and radiation durability—is critical to guiding material selection in future reactor designs. This study revealed quite the opposite of what has been observed for CVD and NITE SiC under irradiation. This could be for a number of reasons and will be explored further with later experimentation.

Looking forward, several research directions can further enhance our understanding and application of SiC in fusion environments. First, more advanced *in situ* experimentation—such as real-time ion irradiation combined with microscopy—could provide dynamic insights into bubble formation and defect migration. Second, expanding machine learning models to include unsupervised or self-supervised approaches may reduce reliance on labeled datasets, allowing faster adoption across new materials systems. Third, the integration of experimental results with multiscale modeling, including molecular dynamics and finite element simulations, could enable predictive frameworks that bridge the gap between atomic-scale interactions and macroscopic material behavior.

References

[1] C. Forsberg, M. Kazimi, E. Moniz, The Future of the Nuclear Fuel Cycle, MIT Energy Initiative Report, 2011.

[2] A. International Atomic Energy, Climate Change and Nuclear Power 2020, IAEA Publications2020.

[3] W.M. Stacey, Fusion: An introduction to the physics and technology of magnetic confinement fusion, John Wiley & Sons2010.

[4] S.J. Zinkle, J.T. Busby, Structural Materials for Fission & Fusion Energy, Materials Today 12(11) (2009) 12-19.

[5] Y. Katoh, L.L. Snead, Silicon carbide and its composites for nuclear applications – Historical overview, Journal of Nuclear Materials 526 (2019) 151849.

[6] S.J. Zinkle, G. Was, Materials challenges in nuclear energy, Acta Materialia 61(3) (2013) 735-758.

[7] M.A. Fütterer, L. Fu, C. Sink, S. de Groot, M. Pouchon, Y.W. Kim, F. Carré, Y. Tachibana, Status of the very high temperature reactor system, Progress in Nuclear Energy 77 (2014) 266-281.

[8] E. Shaber, G. Baccaglini, S. Ball, T. Burchell, B. Corwin, T. Fewell, M. Labar, P. MacDonald, P. Rittenhouse, R. Vollam, Very High Temperature Reactor (VHTR) Survey of Materials Research and Development Needs to Support Early Deployment, Idaho National Lab.(INL), Idaho Falls, ID (United States), 2003.

[9] S.J. Zinkle, G.S. Was, Materials Challenges in Nuclear Energy, Acta Materialia 61(3) (2013) 735-758.

[10] P. Delhaes, Chemical vapor deposition and infiltration processes of carbon materials, Carbon 40(5) (2002) 641-657.

[11] I.M. Low, Advances in the manufacture of ceramic matrix composites using infiltration techniques, Elsevier Science & Technology, Cambridge :, 2014.

[12] J. Hu, C. Liu, F. Ye, L. Cheng, Y. Wei, A review on high-performance SiCf/SiC composites prepared by PIP process, Journal of Materials Research and Technology 33 (2024) 7216-7235.

[13] B. Heidenreich, Carbon fibre reinforced SiC materials based on melt infiltration, 6th International Conference on High Temperature Ceramic Matrix Composites HTCMC 6, 2007.

[14] Y. Katoh, A. Kohyama, T. Nozawa, M. Sato, SiC/SiC composites through transient eutectic-phase route for fusion applications, Journal of Nuclear Materials 329 (2004) 587-591.

[15] A. Gallardo-López, A. Muñoz, J. Martınez-Fernández, A. Domınguez-Rodrıguez, High-temperature compressive creep of liquid phase sintered silicon carbide, Acta materialia 47(7) (1999) 2185-2195.

[16] K.A. Terrani, C. Ang, L.L. Snead, Y. Katoh, Irradiation stability and thermo-mechanical properties of NITE-SiC irradiated to 10 dpa, Journal of Nuclear Materials 499 (2018) 242-247.

[17] J. Chen, P. Jung, Effect of helium on radiation damage in a SiC/C composite, Ceramics International 26(5) (2000) 513-516.

[18] C. Tromas, V. Audurier, S. Leclerc, M.F. Beaufort, A. Declémy, J.F. Barbot, Evolution of mechanical properties of SiC under helium implantation, Journal of Nuclear Materials 373(1) (2008) 142-149.

[19] J.F. Barbot, S. Leclerc, M.L. David, E. Oliviero, R. Montsouka, F. Pailloux, D. Eyidi, M.F. Denanot, M.F. Beaufort, A. Declemy, Helium implantation into 4H-SiC, physica status solidi (a) 206(8) (2009) 1916-1923.

[20] Y. Lin, M. Zhou, X. Tai, H. Li, X. Han, J. Yu, Analytical transmission electron microscopy for emerging advanced materials, Matter 4(7) (2021) 2309-2339.

[21] D.B. Williams, C.B. Carter, D.B. Williams, C.B. Carter, The transmission electron microscope, Springer2009.

[22] M. Shen, G. Li, D. Wu, Y. Yaguchi, J.C. Haley, K.G. Field, D. Morgan, A deep learning based automatic defect analysis framework for In-situ TEM ion irradiations, Computational Materials Science 197 (2021) 110560.

[23] S. Jesse, M. Chi, A. Belianinov, C. Beekman, S. Kalinin, A. Borisevich, A. Lupini, Big data analytics for scanning transmission electron microscopy ptychography, Scientific reports 6(1) (2016) 26348.

[24] E. Ramírez-Aportela, J. Mota, P. Conesa, J.M. Carazo, C.O.S. Sorzano, DeepRes: a new deep-learning-and aspect-based local resolution method for electron-microscopy maps, IUCrJ 6(6) (2019) 1054-1063.

[25] M. Hussain, YOLO-v1 to YOLO-v8, the rise of YOLO and its complementary nature toward digital manufacturing and industrial defect detection, Machines 11(7) (2023) 677.

[26] L. Fu, Rule generation from neural networks, IEEE Transactions on Systems, Man, and Cybernetics 24(8) (1994) 1114-1124.

[27] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, K.-R. Müller, Evaluating the visualization of what a deep neural network has learned, IEEE transactions on neural networks and learning systems 28(11) (2016) 2660-2673.

[28] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, Fiji: an open-source platform for biological-image analysis, Nature methods 9(7) (2012) 676-682.

[29] E. Coto, Dutta, A., Zesserman, A. , VGG Annotation Search and Annotator (VASA), 2020. https://www.robots.ox.ac.uk/~vgg/software/vasa/. 2024).

[30] C. Field, Theiascope: A platform for automated quantitative microscopy image analysis. https://www.theiascientific.com/. 2024).

[31] C.-H. Chen, Y. Zhang, Y. Wang, M.L. Crespillo, C.L. Fontana, J.T. Graham, G. Duscher, S.C. Shannon, W.J. Weber, Dose dependence of helium bubble formation in nano-engineered SiC at 700 °C, Journal of Nuclear Materials 472 (2016) 153-160.

[32] B. Shaver, K. Wheeler, C. Ang, Development of SiC encapsulation for thulium oxide targets as potential nuclear batteries, Journal of Nuclear Materials 608 (2025) 155700.

[33] C. Ang, G. Singh, L. Snead, Y. Katoh, Preliminary study of sintering zero-rupture Fully Ceramic Microencapsulated (FCM) fuel, International Journal of Applied Ceramic Technology 16(5) (2019) 1699-1707.

[34] Ion Beam Materials Laboratory - General Information, Los Alamos National Laboratory (2025).

[35] J.F. Ziegler, M.D. Ziegler, J.P. Biersack, SRIM–The stopping and range of ions in matter (2010), Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 268(11-12) (2010) 1818-1823.

[36] S. Grasso, P. Tatarko, S. Rizzo, H. Porwal, C. Hu, Y. Katoh, M. Salvo, M.J. Reece, M. Ferraris, Joining of β-SiC by spark plasma sintering, Journal of the European Ceramic Society 34(7) (2014) 1681-1686.

[37] T. Koyanagi, Y. Katoh, T. Hinoki, C. Henager, M. Ferraris, S. Grasso, Progress in development of SiC-based joints resistant to neutron irradiation, Journal of the European Ceramic Society 40(4) (2020) 1023-1034.

Appendix

Appendix A: Model Test Images and Statistical Results

This appendix presents the supplemental evaluation conducted using lithium aluminate and palladium nickel test images to further assess the performance and adaptability of the preliminary segmentation models. The images were processed using Theiascope™ software, allowing for consistent quantification of bubble count, mask accuracy, and size across each model configuration. These extended tests provided additional depth to the comparative analysis, revealing how different YOLO architectures generalize to different materials beyond the SiC based dataset that is planned to be evaluated. The findings from this evaluation support the decision to prioritize the Small YOLO model and train over more epochs, as this approach demonstrated the most consistent and promising performance across diverse microstructural conditions.

A.1 Lithium Aluminate Test Image and Data

The lithium aluminate test image was selected to examine how the trained models perform on a material with distinct microstructural features, contrast characteristics, and a lot of smaller features. Although the material differs significantly from SiC, it still exhibits white cavity-like structures suitable for segmentation. By analyzing the models on this sample, the evaluation probed their resilience to noise and their effectiveness in environments with more ambiguous feature boundaries. This served as a stress test to understand model limitations in broader materials applications.

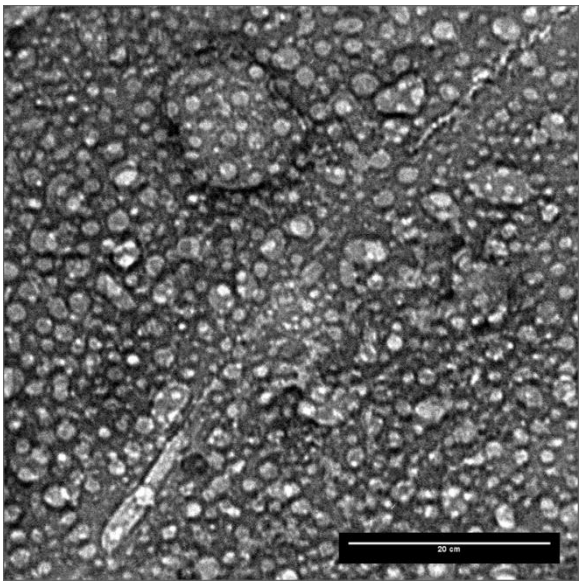

Figure A1. Lithium Aluminate TEM micrograph used for testing which contains 1260 bubbles when manually annotated.

Table A1. Testing Data for Lithium Aluminate for all models trained.

Model	Model Size	Epoch	Count	Accuracy	Size (px)
SAM	N/A	N/A	354		
Model 5	Nano	100	99	0.693	21.8
Model 5	Small	100	105	0.669	22.8
Model 5	Medium	100	154	0.684	21.4
Model 5	Large	100	186	0.694	20.2
Model 5	X-Large	100	124	0.691	22.2
Model 10	Nano	10	118	0.738	22
Model 10	Nano	100	121	0.676	19.6
Model 10	Nano	1000	283	0.732	17.8
Model 10	Small	10	124	0.69	21.9
Model 10	Small	100	191	0.693	19.7

Table A1 Continued

Model 10	Small	1000	403	0.748	17.1
Model 10	Medium	10	134	0.699	21.5
Model 10	Medium	100	160	0.691	20.8
Model 10	Medium	1000	354	0.737	17.4
Model 10	Large	10	58	0.656	21.7
Model 10	Large	100	66	0.661	20.2
Model 10	Large	1000	353	0.741	17.4
Model 10	X-Large	10	322	0.758	18.6
Model 10	X-Large	100	96	0.678	24
Model 10	X-Large	1000	349	0.743	17.6
Model 10 (Unaugmented)	Medium	100	718	0.936	19
Model 20	Nano	100	113	0.688	21.1
Model 20	Small	100	210	0.706	19.9
Model 20	Medium	100	163	0.706	20.3
Model 20	Large	100	196	0.713	19.6
Model 20	X-Large	100	176	0.706	19.8

A.2 Palladium Nickel Test Images and Data

The palladium nickel test images removed most of the additional challenge that comes from higher image complexity and less uniform bubble morphology, as they were selected specifically for the larger bubble feature size. This test helped evaluate the model's ability to

generalize beyond the original dataset without retraining and contributed to the preliminary investigation of how training parameters contribute to the effectiveness of a model.

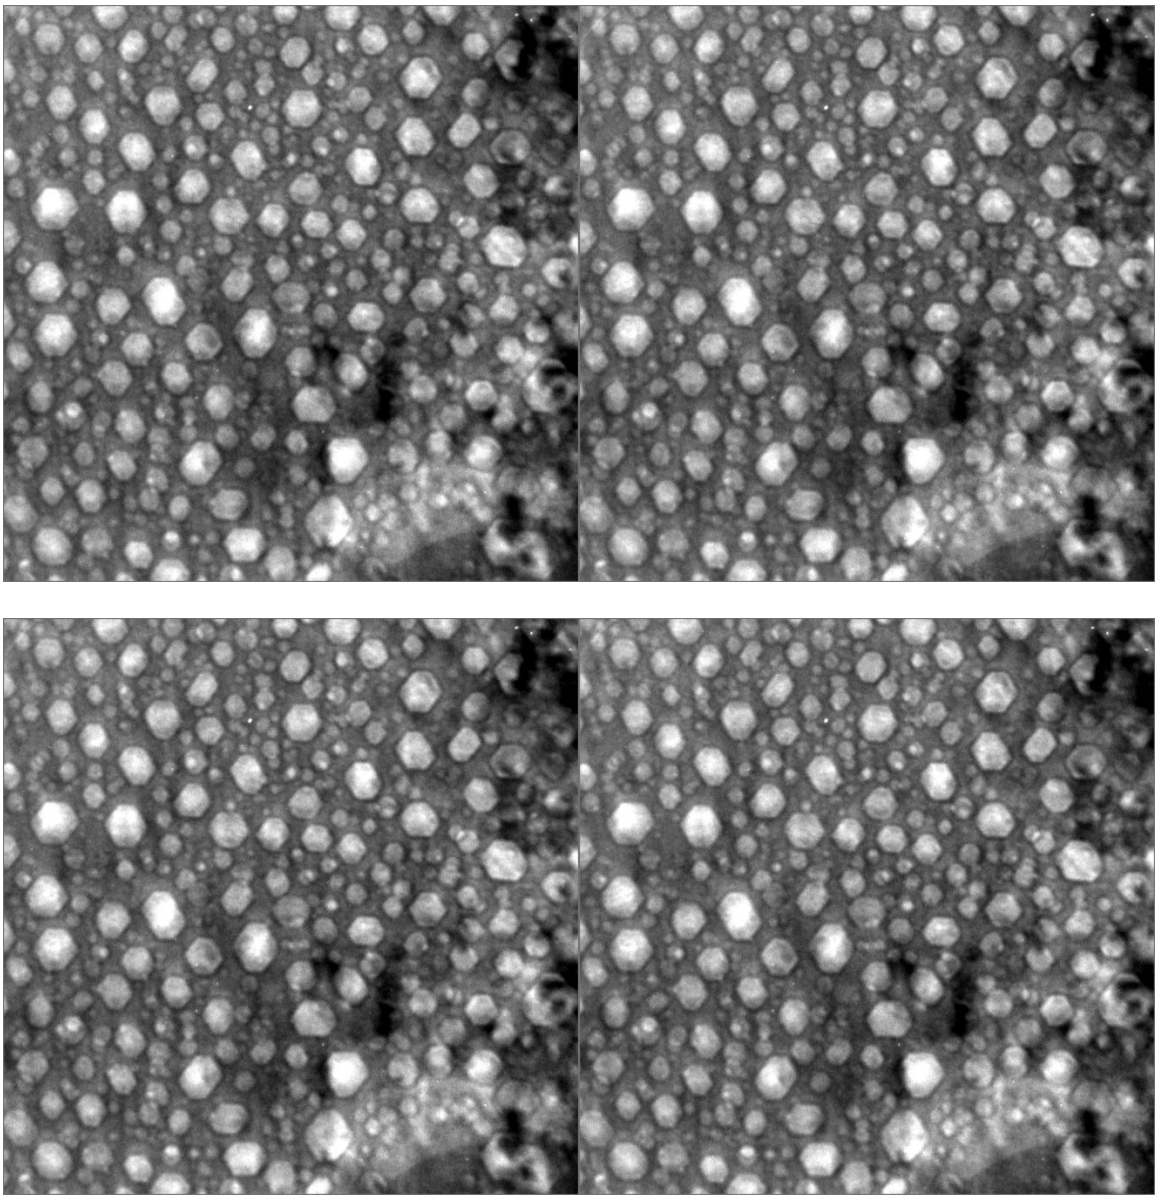

Figure A2. Palladium Nickel TEM micrographs from video frames 16372 through 16376.

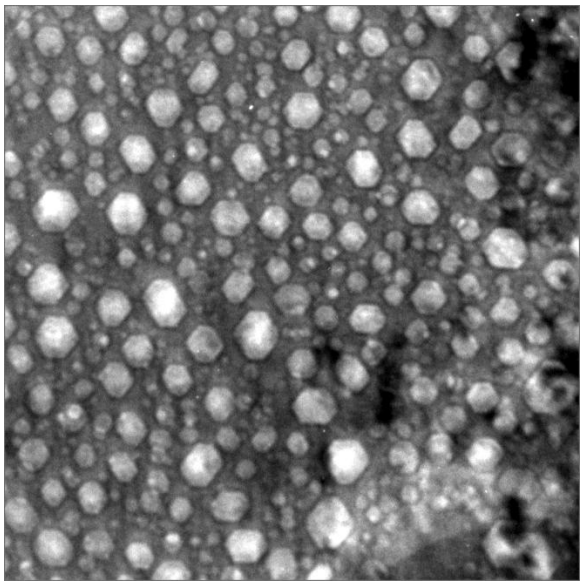

Figure A2 Continued. Palladium Nickel TEM micrographs from video frames 16372 through 16376.

Table A2. Testing Data for Palladium Nickel for all models trained for each frame image used.

Model	Model Size	Epoch	Count	Accuracy	Size (px)
Manual Annotation	N/A	N/A	175	0.8	42.0
			171	0.8	43.6
			185	0.8	40.6
			200	0.8	39.3
			232	0.8	37.9
Model 5	Nano	100	132	0.736	48.3
			128	0.736	48.6
			129	0.733	48.6
			134	0.735	48.1
			139	0.727	47.6

Table A2 Continued

Model 5	Small	100	127	0.743	49.9
			132	0.740	50.1
			133	0.739	48.9
			126	0.748	50.2
			129	0.747	50.1
Model 5	Medium	100	139	0.723	47.7
			134	0.724	47.9
			139	0.723	47.2
			139	0.728	47.5
			140	0.730	47.3
Model 5	Large	100	167	0.773	44.8
			168	0.771	44.4
			169	0.771	44.0
			167	0.773	44.4
			174	0.769	43.8
Model 5	X-Large	100	125	0.759	50.3
			125	0.751	50.8
			126	0.760	50.2
			122	0.769	50.9
			127	0.756	50.6

Table A2 Continued

Model 10	Nano	10	145	0.813	47.3
			150	0.808	46.8
			144	0.814	47.2
			143	0.817	47.5
			151	0.806	46.8
Model 10	Nano	100	153	0.758	45
			150	0.759	45.3
			157	0.752	44.7
			157	0.754	45.3
			153	0.752	46.1
Model 10	Nano	1000	195	0.784	43.6
			194	0.784	43.7
			197	0.782	43.6
			193	0.783	44.1
			194	0.787	43.9
Model 10	Small	10	123	0.711	46.1
			121	0.711	46.0
			127	0.707	44.9
			132	0.703	45.4
			133	0.701	45.3

Table A2 Continued

Model 10	Small	100	144	0.749	49.1
			140	0.749	49.8
			141	0.747	49.2
			140	0.752	49.7
			138	0.754	49.8
Model 10	Small	1000	178	0.789	44.8
			180	0.788	44.4
			181	0.785	44.6
			181	0.785	44.4
			178	0.784	44.9
Model 10	Medium	10	157	0.787	46.9
			153	0.788	47.1
			160	0.783	46.4
			162	0.778	46.4
			156	0.785	47.4
Model 10	Medium	100	175	0.762	43.9
			171	0.764	44.1
			176	0.761	44.5
			175	0.766	44.0
			172	0.767	44.6

Table A2 Continued

Model 10	Medium	1000	209	0.797	43.1
			216	0.791	42.4
			222	0.787	42.1
			228	0.787	41.7
			221	0.789	42.3
Model 10	Large	10	102	0.674	49.6
			98	0.675	49.4
			98	0.672	50.4
			98	0.673	50.8
			104	0.668	48.9
Model 10	Large	100	136	0.726	48.5
			134	0.728	49.3
			134	0.724	48.9
			131	0.731	49.6
			136	0.725	48.8
Model 10	Large	1000	210	0.794	43.3
			211	0.791	42.9
			215	0.786	42.3
			221	0.786	42.0
			217	0.790	42.5

Table A2 Continued

Model 10	X-Large	10	224	0.790	40.5
			221	0.789	40.9
			226	0.785	40.8
			219	0.789	40.8
			224	0.789	40.6
Model 10	X-Large	100	162	0.777	46.6
			159	0.778	47.1
			161	0.777	46.5
			164	0.779	46.6
			163	0.782	46.7
Model 10	X-Large	1000	222	0.783	42.0
			217	0.784	42.1
			218	0.780	42.2
			210	0.788	42.3
			215	0.783	42.2
Model 10 (Unaugmented)	Medium	100	284	0.881	39.0
			275	0.885	38.8
			271	0.881	38.7
			254	0.874	39.9
			273	0.880	38.2

Table A2 Continued

Model 20	Nano	100	165	0.813	50.6
			162	0.815	50.9
			162	0.815	50.9
			162	0.821	50.4
			168	0.815	49.9
Model 20	Small	100	187	0.819	48.1
			188	0.818	47.9
			187	0.820	48.1
			196	0.813	47.6
			193	0.815	47.9
Model 20	Medium	100	178	0.826	48.6
			177	0.823	48.9
			177	0.822	48.8
			178	0.824	48.8
			184	0.820	48.3
Model 20	Large	100	185	0.840	48
			182	0.840	48.2
			182	0.841	48.1
			185	0.840	48.1
			185	0.842	48.1

Table A2 Continued

Model 20	X-Large	100	180	0.815	48.6
			179	0.814	48.8
			178	0.813	48.8
			187	0.807	48.1
			182	0.817	48.6

Appendix B: NITE and CVD SiC Data

This appendix provides supplemental data supporting the broader characterization of the NITE silicon carbide samples analyzed in this study. While the core results focus on hardness trends and segmentation performance, the additional data presented here give further insight into the material's structural and compositional features. Included are the results of X-ray diffraction (XRD) analysis used to evaluate the crystallographic structure of NITE SiC, as well as a detailed breakdown of hardness data for each individual sample. These sections are intended to strengthen the interpretation of mechanical performance and microstructural response by grounding them in well-characterized material baselines.

Appendix B1: X-Ray Diffraction Results

This section summarizes the X-ray diffraction (XRD) analysis performed on the NITE SiC samples. The XRD data provided confirmation of the crystallographic phase and offered insight into structural integrity, grain orientation, and any secondary phases that may influence mechanical or irradiation behavior present in the sample. This also verified the records of what the samples used were made up of given their 7 years worth of aging at room temperature.

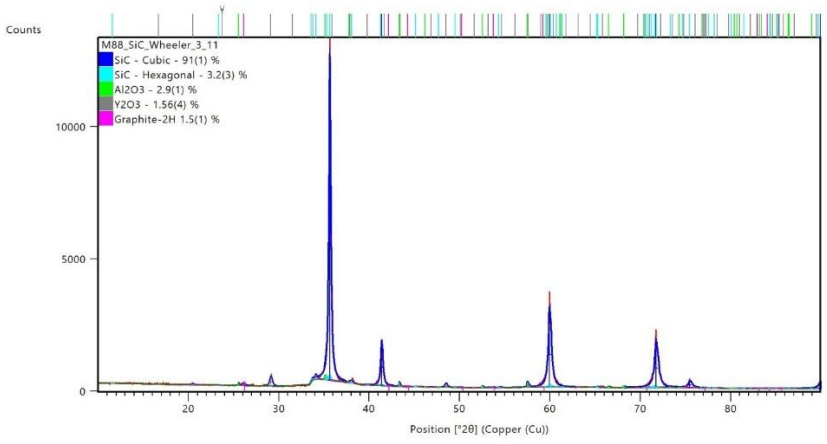

Figure B1. Raw data from XRD analysis of mix 88 NITE SiC.

Appendix B2: Nanoindentation Data

Detailed here are the individual nanoindentation results for all 15 SiC samples. The data includes average hardness values and standard deviations for each implantation condition, offering a clear view of how helium dose influences mechanical consistency and potential outlier behavior at the sample level.

TABLE B1. HARDNESS DATA AND ERROR BROKEN DOWN BY SAMPLE.

Sample	1 x 10^{14} ion/cm^2 (GPa)	Error	1 x 10^{15} ion/cm^2 (GPa)	Error	1 x 10^{16} ion/cm^2 (GPa)	Error
147	39.0851	25.3944	48.1753	5.4840	32.4923	22.5400
150	36.6430	3.0178	26.6316	15.4073	30.0795	4.6365
151	31.5201	2.6170	31.5225	2.30518	39.5282	4.4594
1	32.8785	1.0468	38.5844	1.1830	40.54722	1.1894
2	47.4796	2.9120	44.1792	1.6802	37.9498	1.7713

Vita

Kip Wheeler was born and raised in Edinburgh, IN and went to Columbus North High School. Kip's fascination with science started young and developed into a love for particles and the inner workings of the world around her. To pursue Nuclear Engineering, Kip moved to Knoxville, TN in 2020 to earn a Bachelors Degree in Nuclear Engineering from the University of Tennessee. Kip received her first degree in 2024 and plans to work towards a PhD after the completion of her current Masters work. Kip has a wonderful husband hailing from Taylorsville, IN. They met while attending Columbus North High School and would later have three adorable, chaos-fueled cats of the same Indiana origin to keep them company in Tennessee.