Using Machine Learning for Defect Characterization

Victoria Gerardi
victoria.r.gerardi.civ@army.mil

Antonio Aguirre
antonio.aguirre15.civ@army.mil

Yvan Christophe
yvan.christophe.civ@army.mil

DEVCOM - Armaments Center
Picatinny Arsenal, NJ

ABSTRACT

This paper presents a methodology that is under development to analyze large X-ray image datasets for anomaly and/or defect detection using machine learning techniques. The characterization of anomalies and/or defects can be identified through the performance accuracy of either image classification (supervised learning - convolutional neural networks) or anomaly detection (unsupervised learning - autoencoders) models. Each learning technique has unique hyperparameters and design architectures to aid in creating a robust model to predict against X-ray images of varying orientations, brightness and contrast. This method would be a strong complement to the traditional suite of energetic material/component characterization tests, particularly for melt-pour explosives, performance-related design intent, safety, and/or performance-related defect detection. For safety or performance-related defect detection, the methodology enables baselining defects as a feedback loop in the development of new subscale tests and physics-based models to better understand and predict energetic failure modes, a capability under development at DEVCOM Armament Center called Energetic Defect Characterization (EDC).

ABOUT THE AUTHORS

Victoria Gerardi is an Operations Research Analyst working at Combat Capabilities Development Command – Armaments Center (DEVCOM-AC) in Picatinny Arsenal, NJ since 2018. Victoria earned her Bachelor of Science in Applied Mathematics from the College of William & Mary in 2017 and her Master of Science in Computational Operations Research from the College of William & Mary in 2018. She is the performance analysis direct fire team lead.

Antonio Aguirre works for the US ARMY's Picatinny Arsenal in New Jersey as a mathematician, computer scientist, and data analyst. In his 6-year career with the US ARMY, Mr. Aguirre has most notably: selected to serve as the MATLAB Technical POC for Picatinny Arsenal, Watervliet Arsenal, and Picatinny's Armament University; completed a Radar Systems Certification from Georgia Tech; contributes to, developed, and maintains several tools and models that regularly provide critical insights to senior ARMY and program leadership. Prior to Mr. Aguirre's service with the ARMY, he lived in New York City where he earned a Bachelor of Science in Applied Mathematics with the distinction Magna Cum Laude. While in college Mr. Aguirre was selected to work with world class scientist in the field of Remote Sensing at NASA-Goddard Space Flight Center, NASA-Glenn Research Center, and twice at Brookhaven National Laboratory.

Yvan Christophe is a General Engineer working at Combat Capabilities Development Command – Armaments Center (DEVCOM-AC) in Picatinny Arsenal. He obtained is master's degree from Stevens Institute of Technology and majored in Mechanical Engineering while concentrating in mechanical design, manufacturing, and robotics. Presently, Mr. Christophe works within the System Analysis Division.

Using Machine Learning for Defect Characterization

Victoria Gerardi
victoria.r.gerardi.civ@army.mil

Antonio Aguirre
antonio.aguirre15.civ@army.mil

Yvan Christophe
yvan.christophe.civ@army.mil

DEVCOM - Armaments Center
Picatinny Arsenal, NJ

INTRODUCTION

Current and future energetic requirements in artillery munitions are exceeding legacy gun and barrel designs as well as flight environments. This means that non-critical defects of today may become critically defective in the future. Defective artillery has caused catastrophic failures at gun launch, injuries and fatalities to personnel, and damage to platforms (Ismay; Kumar; Singh). Preventing and mitigating these potentially critical defects is a priority for the U.S. Army. While there are processes in place to detect defects, these methods are time intensive, expensive, and outdated. It is time to establish the detect detection methodology of the future. Given the impact of this task, the U.S. Army Combat Capabilities Development Command - Armaments Center is working to develop and demonstrate modern defect detection capabilities.

The Capability Development Effort - Energetic Defect Characterization (CDE - EDC) is a new capability development program funded by U.S. Army Combat Capabilities Development Command Armament Center at Picatinny Arsenal. This effort aims to develop capabilities to enable experimental and computational evaluation & prediction of energetics with defects, especially for enhanced munitions (LRPF). In addition to the working group located at Picatinny Arsenal, there is also international cooperation with the NATO Working Groups and The Technical Cooperation Program (TTCP). Additionally, analysts are also collaborating with cadets at the United States Military Academy (USMA) at West Point. CDE – EDC is comprised of three core sub-groups: Experimental, Data Analysis, and Computational, all of which are located at DEVCOM-AC. The relation between the three sub-groups is seen in Figure 1. This report focuses on the work completed by the Data Analysis sub-group. With the combined technical expertise of these three sub-groups, DEVCOM-AC is striving to provide updated capabilities and guidance for defect detection, characterization, and mitigation.

The Experimental Testing Sub-Group is in the Energetics, Warheads, and Manufacturing Technology Directorate and is tasked with identifying and quantifying the fundamental physical and chemical mechanism(s) that may drive unintended ignition of defective energetics. There are a variety of physical tests being completed by this sub-group to collect this data. Part of this task incudes improving realistic defect replication for the physical testing.

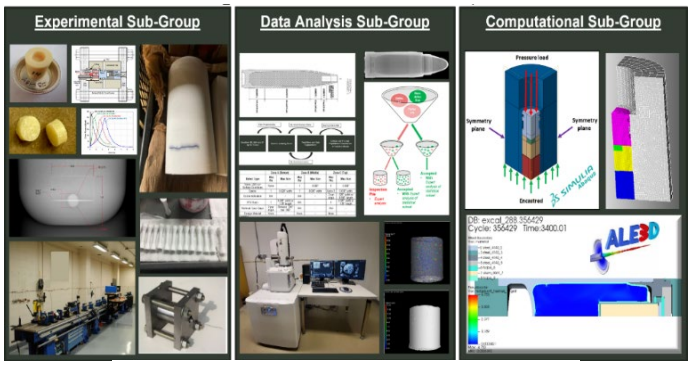

Figure 1: EDC Sub-Groups

The Data Analysis Sub-Group is in the Systems Analysis Division and is tasked with using machine learning to flag and characterize anomalistic and defective inspection images. The knowledge gained from using these models will aid both the experimental and computational sub-groups in their tasks.

The Computational Sub-Group is in the Energetics, Warheads, and Manufacturing Technology Directorate and is tasked with modeling defects in energetics to assess whether premature ignition is likely or not. This sub-group simulates gun launches given energetic parameters to determine effects of defects. This sub-group uses models such as Abaqus and STAR-CCM+. The computational sub-group compares model outputs to the physical test outputs from the Experimental sub-group. As such, it relies heavily on the experimental sub-group and vice versa.

The deliverable of this effort will be a self-sufficient stand-alone predictive tool that can be utilized by the radiologists examining manufactured munitions. This report will focus on the machine learning portion of this effort.

CURRENT METHODOLOGY

The current defect detection methodologies are time intensive, expensive, and outdated. While the specific aspects of a methodology depend on project/munition requirements, the general process involves multiple radiologists examining X-ray/CT scans of manufactured munitions. There are many types of defects: gas porosity, gaps, voids, cracks, etc. The size, type, and location of these defects all come together and may or may not result in a critical defect. However, recreating critical defects in these specific types is a difficult challenge, and is one that the current effort is also investigating experimentally (Lin).

By leveraging machine learning methods, the goal is to be able to determine whether a critical defect is present as well as gaining a clearer understanding of defect characterization, and, indirectly, their effects. Figure 2 illustrates our "crawl, walk, run" approach for accomplishing these goals. To "crawl" we conducted a standard literature review focusing on industry standards for manufacturing defect detection (Ferguson MK). This yielded convolutional network architectures, as well as transfer learning models that served as a valid starting point for this study. To "walk" we are implementing machine learning models with increasing complexity. Building up from a simple binary classification model, towards a multi-class defect detection and characterization. To "run" we are maturing and consolidating our developed model and tools into a deployable form factor which will be extendable to new defects and munitions.

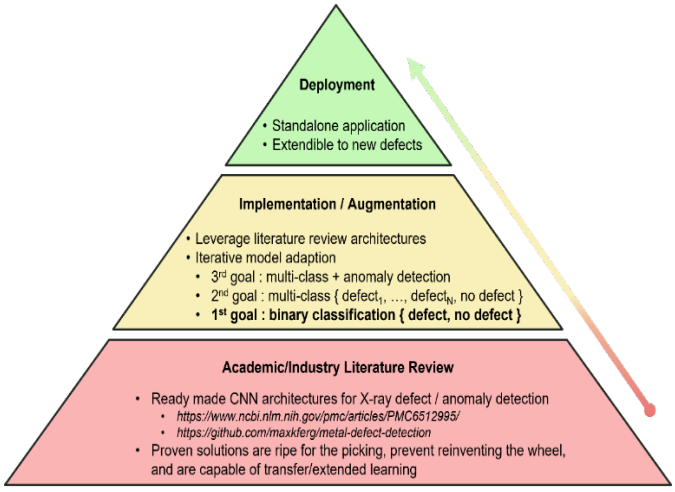

Figure 2: Crawl, Walk, Run

COMPUTATIONAL TOOLS

The Python Programming Language was used for both data preparation as well as model building and training. Standard (accepted) python packages for machine learning were used for this effort, such as pytorch, tensorflow, numpy, pandas, opencv, and labelme.

Analysts initially began work on this effort using the Artificial Intelligence Research Infrastructure (AIRI), which is a local, unclassified high performance computing cluster located at Picatinny Arsenal. They also made use of the Department of Defense High Performance Computing-Modernization Program (DOD HPC-MP) on the Onyx cluster. After methodical experimentation with image analysis and model training on HPC resources, analysts realized that the computational power of the HPCs were not necessary for the current munition data set. Analysts were able to efficiently transfer the model workflow over to their laptops without any loss in performance.

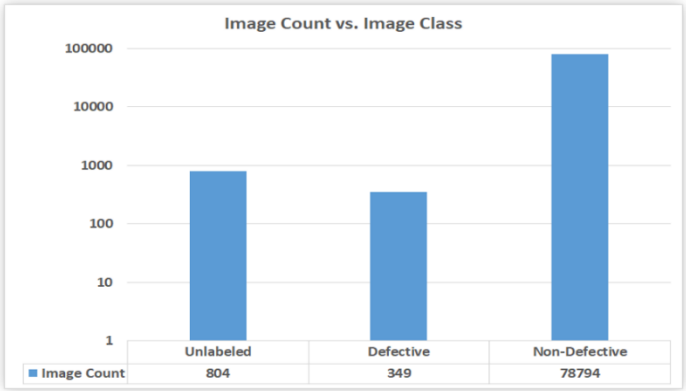

Figure 3: Type-2 Image Classes

INPUT DATA

Exploratory Data Analysis

The munition focused on in this report will be referred to as Type 2. The X-ray images for this munition were captured from 2020-2021. The X-ray images were captured for one aspect angle.

The dataset of all images amounted to 680 GB, with 79,947 munitions captured. Out of these images, 349 munitions were classified as defective by radiologists, representing 0.44% of all munitions. This is an extreme imbalance between the two classes of non-defective vs. defective. This class imbalance will have to be considered when model building. In addition, 804 images were unlabeled. These images were removed from the data set. Figure 3 depicts the split among the image classes.

Data Quality

The requirements document divides the munition into four vertical zones as depicted in Figure 4. These zones are referred to as zones A, B, C, and D. Zone A is at the base of the munition and zone D is close to the nose. According to the requirements document, each zone has a different threshold regarding what constitutes a defect or not. Newer munitions are undergoing extremely high altitudes, high temperatures, set back and vibration shock as well as impact forces. Due to the physics of a gun launch, the base of the munition (zone A) is subjected to higher pressure and temperature than other zones of the munition. Therefore, there are different defect requirements for each munition zone.

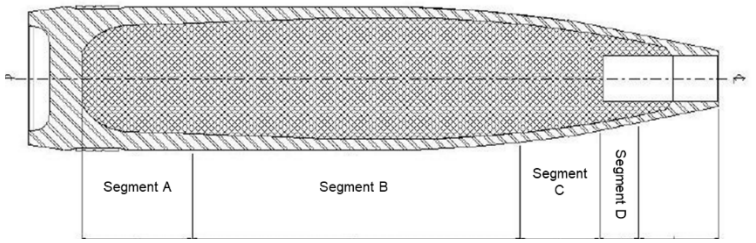

Figure 4: Munition Zone Segmentation

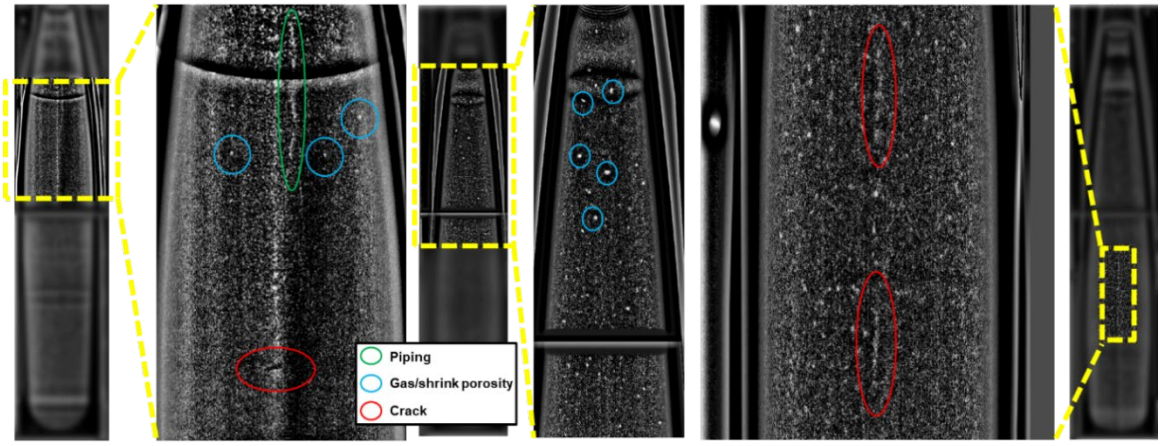

Figure 5: Munition Anomalies

During a manual image inspection of Type-2 munition X-rays, anomalies were found and confirmed. These anomalies were not confirmed to be defects according to the munition requirements. Figure 5 depicts three different munitions, all randomly selected, with their anomalies pointed out. These munitions display signs of piping, gas porosity, and cracks. However, none of these munitions were labelled as defective, due to their location within a specific munition zone, combined with their shape and size. Because the location is so important to the defect determination, images will be segmented into individual zones before model training. If images were not segmented by zone, before being used in the machine learning model, then the model would have a hard time determining which munitions are defective or not. This is attributed to the presence of non-defective anomalies in a zone that are defined as defects by the requirements document in another zone and vice versa.

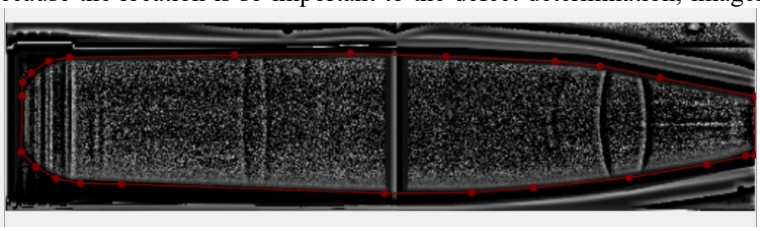

Figure 6: Noisy Background

Additionally, the background removal of each munition was explored to improve model performance. The analysts hypothesized that the background represents noisy and distracting data that is outside the munition region of interest and is not relevant to munition defects or anomalies. Analysts predicted that including the munition background would

make training more difficult for the neural network. Figure 6 is an example of a noisy background that would making model training more difficult. In this figure, everything outside of the red border would be considered noise. However, this ended up not being the case for every munition zone. For Zone A, the background was necessary to aid model training. This will be covered in more detail in the Results section of this report. For Zone B, the background was removed.

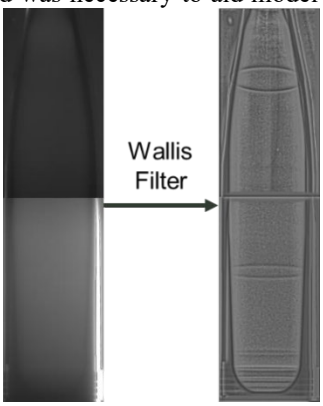

Figure 7: Unfiltered Image and Wallis Filter Effect

Data Processing

Each image received is a 16-bit image, meaning that there are 65536 possible gray values for each pixel. To make anomalies and defects stand out in this image when examining the X-rays, the radiologists suggested for the analysts to make use of the Wallis filter. The Wallis filter is a locally adaptive contrast enhancement filter that raises the contrast in low contrast regions and lowers the contrast in those sections that are too high (Navard). An example of an unfiltered (raw) X-ray image of a munition and the effect of the Wallis filter is captured in Figure 7. The Wallis filter function is defined below in Equation 1. The Wallis filter was not used to process images feeding into the machine learning model but was used by analysts during munition masking.

$$pixel(i,j) = \alpha * \mu + (1-\alpha) * m(i,j) + \big(pixel(i,j) - m(i,j)\big) * \frac{\sigma}{\sigma + limit} \quad (1)$$

α: Brightness factor
μ: Target mean
σ: Target standard deviation
$m(i,j)$: local mean
$limit$: Gain factor

The next step of data processing is the munition masking. This was done using the python package labelme and a U-net. Labelme is Python package which is a graphical image annotation tool and was inspired by a project at MIT (B. C. Russell). Labelme allows for image annotation, image flag annotation, video annotation and more. For this application, analysts used the Wallis filter and labelme to create a dataset of labeled munitions, where the body of the munition was delineated and saved. An instance of the labeling is seen in Figure 8. The body of the munition is "cut out" and marked with the munition name.

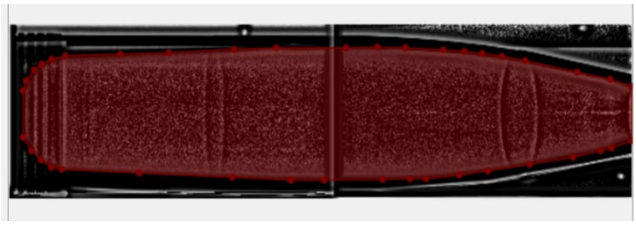

Figure 8: Labelme Functionality

While the filtered images were used by the analysts to clearly see the border between munition and background, the input data to the U-net are the unfiltered (raw) munition images, as along with the .json file of the munition masking. A U-net is a convolutional network architecture that is used for image segmentation. A U-net, also called a "fully convolutional network" and was developed by the University of Freiburg in 2015 (Ronneberger). The architecture of the U-net is depicted in Figure 9. The images are read into the network and are propagated through all possible paths. The output of the U-net is a segmentation, or a mask of the image. The first portion of the U-net is a contraction, consisting of typical convolutional layers with a nonlinear activation function (ReLu) and max pooling layers. The second part of the architecture is the expansion. This portion makes use of up-sampling followed by a convolutional layer, known as "up-convolution" to get back to the original image size with labelled segmentation. The output of the U-net is a predicted mask for the munition border, which is used for both background removal and zone segmentation.

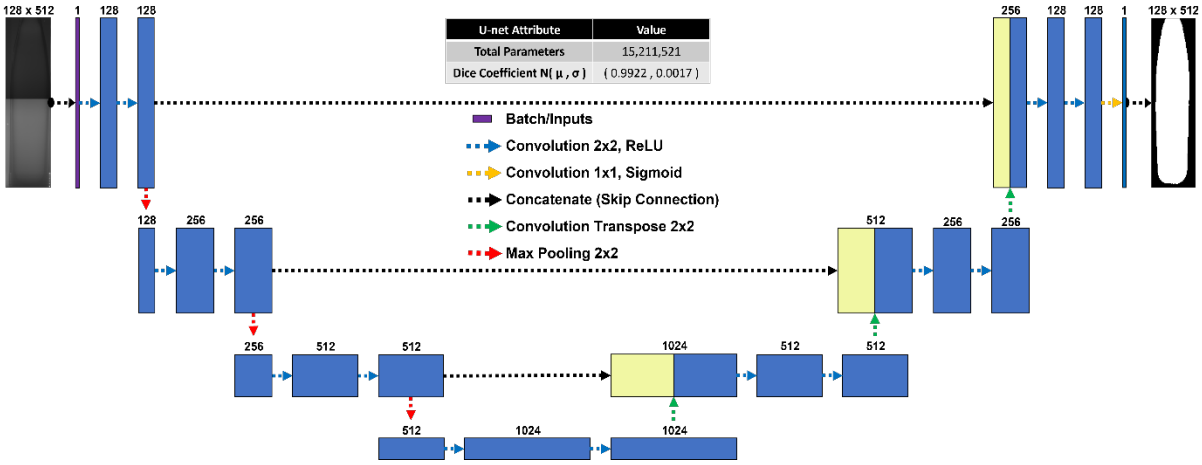

Figure 9: U-net architecture

One benefit to the U-net is that it can train on a small, annotated input data set. The analyst originally manually labelled 100 images. During training iterations, this eventually grew to 950 munition X-ray images. Note that this is less than 1.2% of all munition images. Before being fed into the network, the images were resized from their original 992 x 3878 pixel dimensions to 128 x 512. Using smaller images with the U-net allowed analysts to train a larger AL model, capable of learning more features within the data. It also reduces disk pace and training time requirements. Additionally, the smaller image pixel dimensions are in powers of 2. This removes complications from the U-net layer sizes. The annotated data set has approximately 50% non-defective images and 50% defective images. Having this higher split between non-defective and defective munitions allows for more robust model training to account for variation of munition backgrounds in both non-defective and defective images. The U-net trained for a maximum of 200 epochs with the Adam optimizer. However, an early stopping condition was implemented. If the validation loss worsened by at least 0.01 relative to the last best recorded measure during training over 20 epochs, then the training ended, training ended. The validation loss function for this network was binary cross entropy loss combined with a sigmoid, resulting in a logit output.

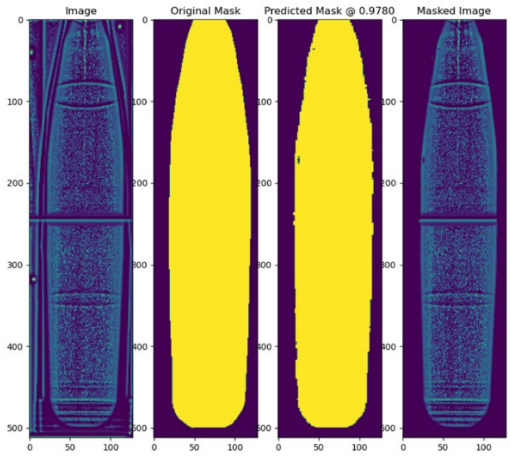

Figure 10: U-net Munition Masking Example

The metric used to test performance of this network is the Dice coefficient, which is used to measure the similarity between two samples, often for image segmentation. An example of the background removal is seen in Figure 10. The U-net performance for the munition masking had a mean Dice coefficient of 99.21%, with a standard deviation of 0.19%. Any holes that are present in the predicted mask are fixed by using a convex hull around the prediction mask. The Dice coefficient function is defined below in Equation 2.

$$Dice = \frac{2 * TP}{2 * TP + FP + FN} \tag{2}$$

TP: True Positive
FP: False Positive
FN: False Negative

The next step of data processing is splitting the munition into its respective zones. This is an automated process and is based off of the munition masking outputted from the U-net. Because there is minimal variation in the height of the predicted masks of the munitions, each zone segment of the munition is a constant pixel height. Thus, each munition image is divided into the four munition segments according to the requirements document. An example is depicted in Figure 11. The full data processing workflow is illustrated in Figure 12.

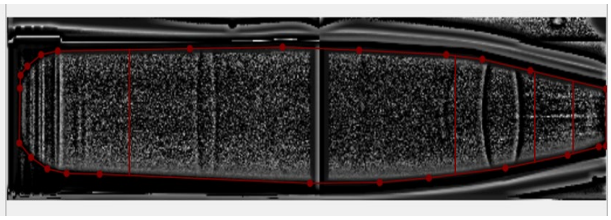

Figure 11: Zone Division

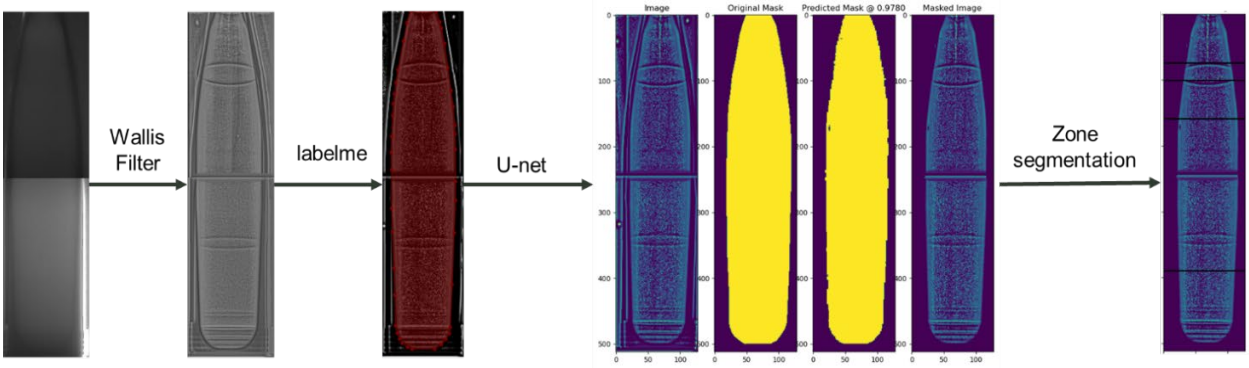

Figure 12: Data Processing Workflow

MODEL BUILDING

The model building workflow will be composed of three portions, depicted in Figure 13. Each layer acts as a funnel, separating defective and non-defective images.

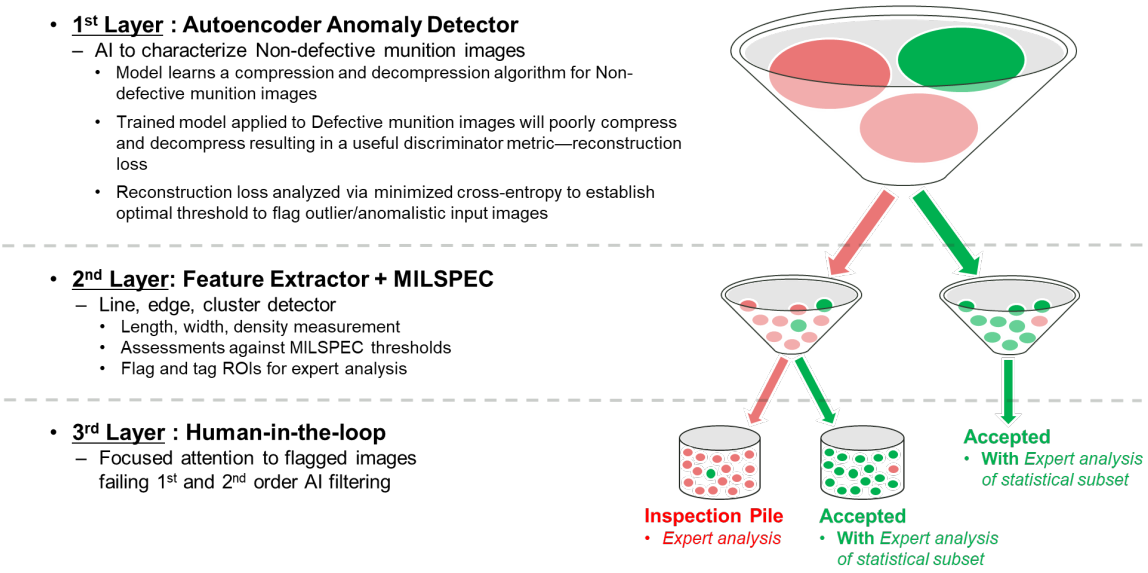

Figure 13: Model Workflow

Using Variational Autoencoders

The top layer of the funnel is a variational autoencoder. The inputs to this neural network are the zonal segments of the filtered munition image. An autoencoder is a type of neural network that can learn a latent representation of the input space. This type of architecture is often referred to a "bow-tie" shape because it has two distinct portions to the network. The first half of the network is called the encoder; it decreases the input shape into what is known as the

latent space, which learns the most significant patterns in the data. The encoder learns the normal distribution (represented by μ and σ) of the input space. Gaussian noise is then sampled from this distribution. The second half of the network, the decoder, is symmetrical to the encoder and recomposes the signal from the Gaussian noise in the latent space. The loss function represents the difference between the input and the output, known as the reconstruction loss. The structure of our network is in Figure 14. There are 2 layers each in the encoder and decoder, with 700 fully connected nodes in the first layer and 300 fully connected nodes in the second layer.

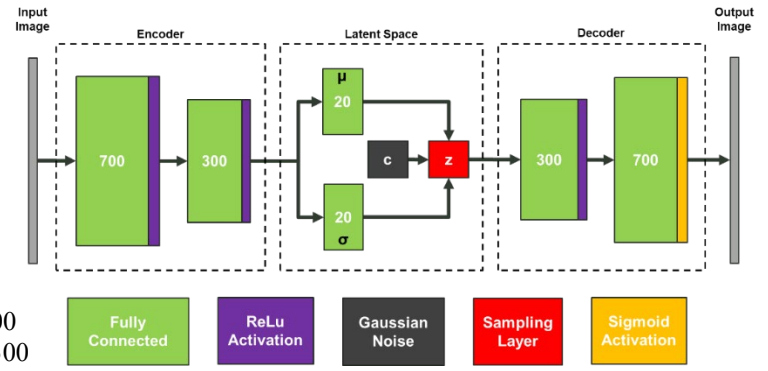

Figure 14: VAE Structure

For this effort, the variational autoencoder is trained on a dataset with only non-defective images. By training the autoencoder only on the non-defective images, the result is a machine learning model that has learned a compression and decompression algorithm for non-defective munition X-rays. Then, when it is tested on a dataset with a mix of both non-defective and defective images, the defective image has a higher reconstruction loss, thus setting them apart from images without defects. This reconstruction loss error for each pixel defined as a linear combination between mean squared error and Kullback–Leibler divergence. These functions are defined in Equation 3.

An example of how the reconstruction loss is illustrated in Figure 15. Blue segments are non-defective samples and defective samples are red. Because the autoencoder has been trained on non-defective images, they are expected to have a very low reconstruction loss. Likewise, the defective images often have higher reconstruction loss. However, there is still the opportunity for misclassification. A threshold value is set for the reconstruction loss; any reconstruction loss value below the threshold is classified as non-defective, and anything above is classified as defective. There can be defective images with reconstruction loss below the anomaly threshold and vice versa. Further investigation will occur for defective images falling into the "high-risk" area, i.e. defective munitions getting classified as non-defective.

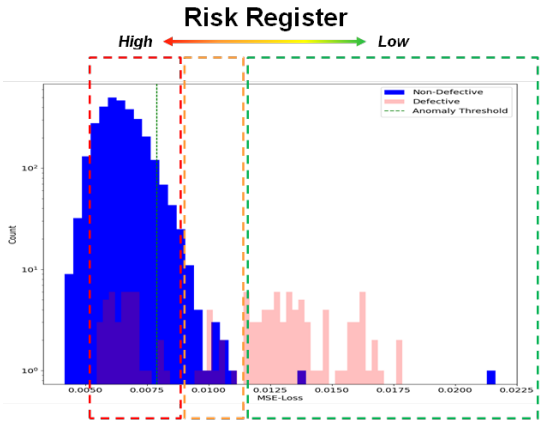

Figure 15: VAE Results

$$MSE(y_{pred}, y_{true}) = \frac{\sum(y_{true} - y_{pred})^2}{n}$$

$$KLD(y_{pred}, y_{true}) = y_{true} \cdot \log\frac{y_{true}}{y_{pred}} = y_{true} \cdot (\log y_{true} - \log y\_pred) \quad (3)$$

$$Loss = 1X10^{-5}MSE + 1X10^{5}KLD$$

Feature Extraction

The second layer of this plan is a feature extractor combined with the requirements document. A line, edge, and cluster detector, combined with measurements of length, width, and density. With this, the anomalies in each zone can be measured against the munition requirement thresholds to determine if the anomalies should be classified as defective. Along with this, the regions of interest can be flagged for further analysis. At this point, the output of the second layer is an aggregation of the results from its zonal segments.

Human-in-the-Loop

The last layer of this plan is the radiologist, i.e., the human-in-the-loop. This human will provide focused attention to flagged images failing the 1st and 2nd layer of the model building workflow. This will allow for robust model tuning as training continues.

RESULTS

The results of the variational autoencoder are based on a test data set for each zone. Because each zone has different requirements for defect characterization; as well as manufacturing differences, there is not a consistent number of defects for each zone. The breakdown of defects per zone is listed in Table 1. For this report, both Zona A and Zone B has been trained using the variational autoencoder, with promising results. Moving forward, Zone C will be examined as well. Due to the limited number of defects present in Zone D, there are currently challenges to producing a useful model capable of predicting defects in Zone D. For Zone A, the test data set was 5119 images, and for Zone B, the test data set consisted of 221 images. The autoencoder in its current state can accurately identify defects current within both Zones A and B of a munition.

Table 1: Number of Defects per Zone

Zone	Number of Defects
A	123
B	44
C	75
D	2

Zone A Results

The confusion matrix for model performance on Zone A is depicted in Figure 16. The model was able to identify 100% of defective samples of Zone A. It also identified 99.94%of non-defective samples in Zone A. This model has a precision score of 0.98, and a recall score of 1.00. The reconstruction loss histogram shows a relatively clear divide between defective and non-defective images. There are three non-defective samples which had a reconstruction loss above the threshold for anomalies. During training, analysts concluded that including the background of the munition in the training set dramatically increased the test accuracy. Without the background included the test accuracy on defective images was approximately 70%. Including the background information allows for discriminating information in the X-ray that aids the model in distinguishing between defects within the munition energetic and anomalies that are present on the outer munition body.

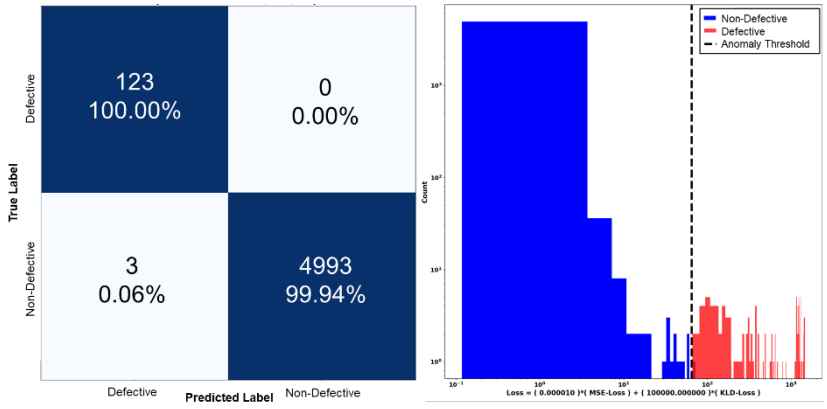

Figure 16: Zone A Performance

Zone B Results

The confusion matrix for model performance on Zone B is depicted in Figure 17. The model performed was able to identify 100% of defective samples of Zone B. It was able to identify 98.99% non-defective samples in Zone B. This model has a precision score of 0.47, and a recall score of 1.00. While this model performed very well at detecting defects, there was a higher percentage of non-defective images

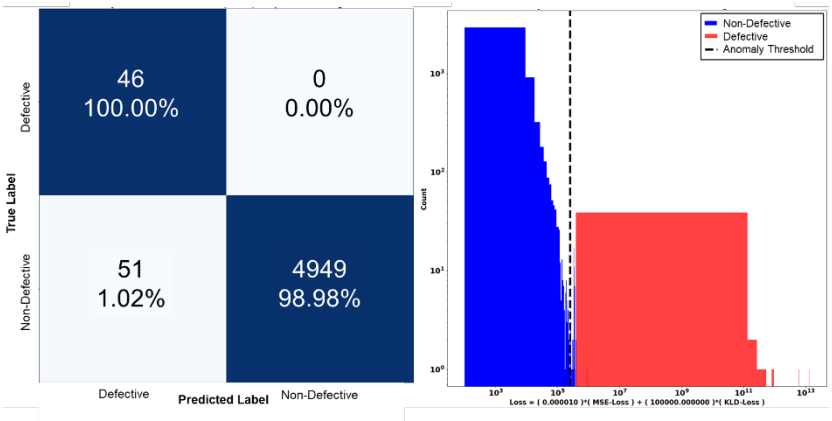

Figure 17: Zone B Performance

classified as defective than the Zone A model.

Zone C Results

The confusion matrix for model performance on Zone C is depicted in Figure 18. The model performed was able to identify 99.43% of defective samples of Zone C. It was able to identify 76.00% non-defective samples in Zone C. This model has a precision score of 0.13, and a recall score of 0.99. While this model performed well at detecting defects, there was a much higher percentage of non-defective images classified as defective than the Zone A or Zone B model. The analysts are currently looking into tuning the hyperparameters of this model using Bayesian optimization. Additionally, they are investigating this data set with expert radiologists on the team.

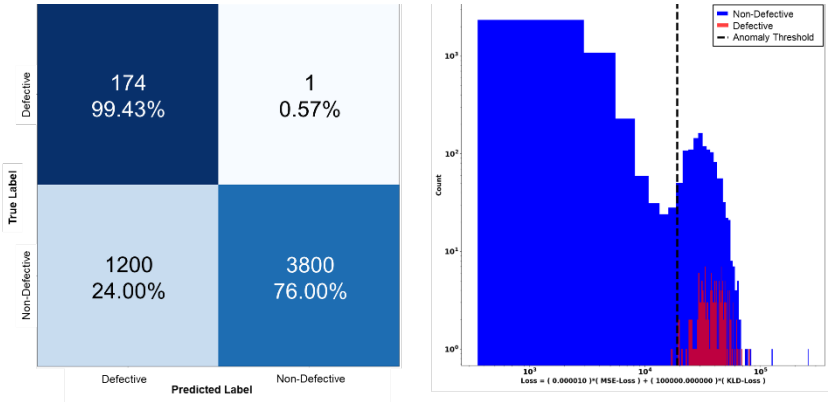

Figure 18: Zone C Performance

SUMMARY

To summarize, the CDE-EDC is steadily working towards our goal of improving safety for the soldier. With outdated requirements, outdated defect detection schemes, and ambiguous defect characterizations, it's time to improve the U.S. Army's defect detection capabilities.

Focusing on munition Type-2, a Wallis filter was used to increase the contrast of the X-ray images and make the munition body more apparent within the X-ray distinct. With the munition body visible, an annotation tool was used to create mask images that delineate the munition body from the image background. Then, a U-net was trained on the X-ray and mask images, to a high-degree of performance, and used to separate the background from the munition, so that the munition could be segmented into its corresponding zones accurately. The munition data was used to train the machine learning model. This U-net worked very well, with a 99.21% Dice rate. Munition X-rays were then split according to their zones, as defined by requirements, to train the model according to different defect thresholds for each zone.

These zonal segments were used as inputs to a variational autoencoder which was trained on a training set of non-defective images. Using it to predict on a test data set made up of both non-defective and defective images proved effective with a high accuracy rate.

FUTURE WORK

There are many directions to take this work in moving forwards. Beginning with our Data Analysis sub-group, we can improve the variational autoencoder used to detect defects in Zone B. The current threshold minimizes a linear combination of pixel-wise MSE and Kullback–Leibler divergence. However, there may be a way to increase accuracy on non-defective samples in Zone B. Additionally, this model methodology still needs to be implemented and trained on Zone C of the munition.

Additionally, this model is currently a binary classifier. It can be expanded to a multi-class classifier and determine defect types (such as gas porosity, crack, piping, etc.). Characterizing the distinct defect types will feed directly into the experimental and computational sub-groups and their task of recreating and simulating defects. Furthermore, this model can also be expanded/adapted for CT scan data. While X-ray images are 2-dimensional, CT scans are 3-dimensional. This volume dimension could have valuable information regarding defect classification and characterization. Lastly, this model can be trained on other munitions.

The things our counterparts can improve on would be capturing data with machine learning in mind. This involves the radiologists recording their findings in a consistent and cohesive manner, as well as capturing meta data that is not

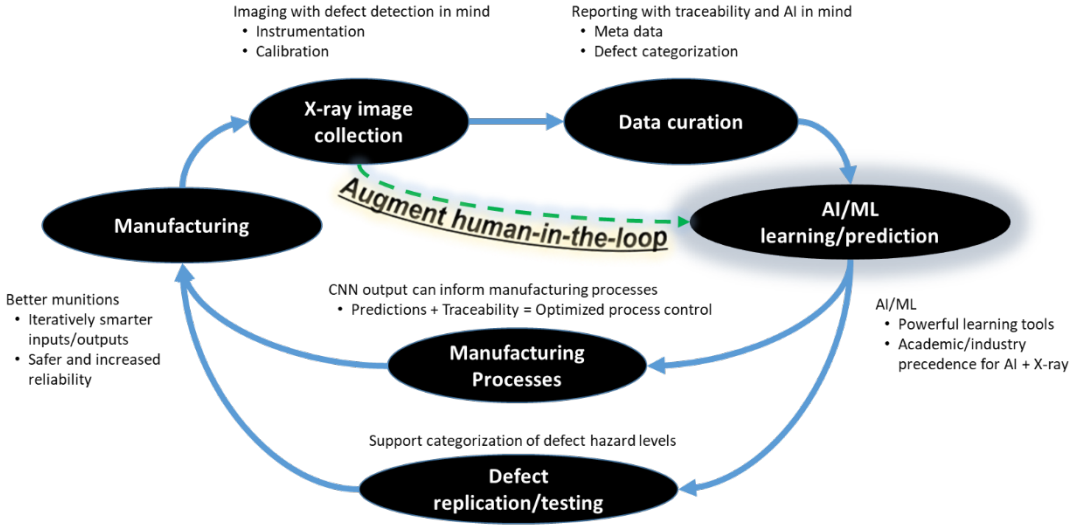

Figure 19: Where we fit into the big picture

currently available to us, such as instrumentation and calibration. It's possible some of this meta data could aid the model prediction and robustness. If data is collected and curated with traceability in mind, it reduces time spent cleaning, as well as tracking down necessary data.

The work chronicled in this report can have wide reaching impacts on the traditional munition manufacturing process. There is the opportunity to augment the human-in-the-loop (the radiologist) in the defect detection cycle. While the human in this process should never be eliminated, having this machine learning model fully integrated at the X-ray site would be valuable to the highly trained technician; either by flagging the cases they should pay the most attention to (munitions that might be just barely defective, i.e., the "difficult" ones), or by tagging munitions with several critical defects. Additionally, this machine learning model will also help fill knowledge gaps that currently exist in the Experimental and Computational sub-groups. Enhancing the understanding of defect characterization will allow these sub-group to hone in on particularly important defects. This work of replicating and testing/simulating defects will in turn lead to updating requirement documents based on defect thresholds that are traceable to specific tests and analysis; improving manufacturing processes to minimize defects; and mitigating the effects of munitions defects. These impacts of this growing capability are highlighted in Figure 19.

Deployable Capability

This work has multiple opportunities of deployment. This first is the deployment of the Wallis filter. This has been implemented into a desktop executable and is shown in Figure 20. This tool allows for images to be dragged and dropped, and all parameters can be adjusted for the individual application.

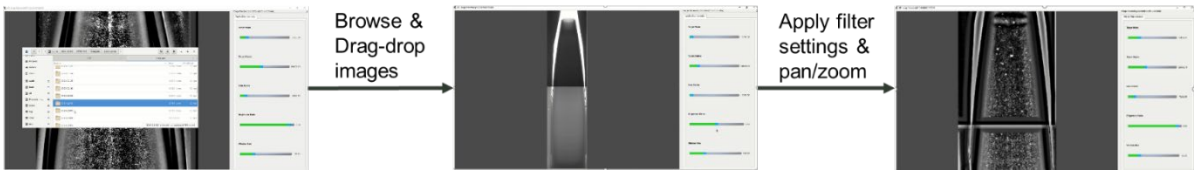

Figure 20: Deployable Wallis Filter Tool

Additionally, there is the possibility for the entire model workflow to be implemented at the X-ray machine with the radiologist. Because the model does not require HPC resources to train, it can be implemented on a laptop or small computational device and rigged up into the current radiologist workflow to improve their performance.

ABBREVIATIONS AND ACRONYMS

AI:	Artificial Intelligence
CDE-EDC:	Capability Development Effort - Energetic Defect Characterization
CNN:	Convolutional Neural Networks
DEVCOM-AC:	Combat Capabilities Development Command – Armaments Center
DOD HPCMP:	Department of Defense Higher Performance Computing Modernization Program
EDC:	Energetic Defect Characterization
HPC:	High Performance Computer
LRPF:	Long Range Precision Fire
ML:	Machine Learning
MSE:	Mean Squared Error
TTCP:	The Technical Cooperation Program
USMA:	United States Military Academy
VAE:	Variational Autoencoder

REFERENCES

B. C. Russell, A. Torralba, K. P. Murphy, W. T. Freeman. "LabelMe: a database and web-based tool for image annotation." International Journal of Computer Vision (May, 2008): pages 157-173, Volume 77, Numbers 1-3.

Ferguson MK, Ronay A, Lee YT, Law KH. "Detection and Segmentation of Manufacturing Defects with Convolutional Neural Networks and Transfer Learning." Ed. NIST. Smart Sustain Manuf Syst. , 2018.

Ismay, John, and Thomas Gibbons-Neff. "Artillery Is Breaking in Ukraine. It’s Becoming a Problem for the Pentagon." New York Times 25 November 2022. Article.

Kumar, Dinesh. "Defective Ammunition: Whose Life Is It Anyway?" The Sunday Guardian Live 18 May 2019. Article.

Lin, Zhong, et al. "Defect engineering of two-dimensional transition metal." Ed. 2d Materials. 2016.

Navard, Andy, et al. "Wallis Normalization Image Enhancement Tactical Support System (TESS (3)) Documentation." DTIC (1992). apps.dtic.mil/sti/pdfs/ADA248301.pdf.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." International Conference on Medical image computing and computer-assisted intervention. Ed. Springer. 2015.

Singh, Rahul. "Army's new US-made M777 howitzer in trouble, barrel explodes at Pokhran range." 12 September 2017. Hindustan Times. 2022.