

**TOWARDS AUTOMATED AND EXPLORATIVE
CHARACTERIZATION OF NANO-ENERGETIC MATERIAL
RESPONSE TO DIRECTED ENERGY**

A Thesis presented to
the Faculty of the Graduate School
at the University of Missouri

In Partial Fulfillment
of the Requirements for the Degree
Master of Science

by
TIMOTHY BRADLEY GAINES
Dr. Grant J. Scott, Thesis Supervisor
December 2023

The undersigned, appointed by the Dean of the Graduate School, have examined the thesis entitled:

TOWARDS AUTOMATED AND EXPLORATIVE
CHARACTERIZATION OF NANO-ENERGETIC MATERIAL
RESPONSE TO DIRECTED ENERGY

presented by Timothy Bradley Gaines,
a candidate for the degree of Master of Science and hereby certify that, in their opinion, it is worthy of acceptance.

Dr. Grant J. Scott

Dr. James M. Keller

Dr. Matthew R. Maschmann

ACKNOWLEDGMENTS

Thank you to my thesis advisor, Dr. Grant Scott for his support, guidance and experience I have gained through this opportunity. Additionally, thank you to the members of my committee: Dr. Hurt, Dr. James Keller and Dr. Matthew Maschmann for their technical expertise and advice. I would also like to thank my co-workers and fellow graduate students for their assistance and help over the years as well as their friendship. I want to especially acknowledge Camden Boyle, Trevor Bajkowski, and David Huangal.

Lastly, I would like to thank my wife Elizabeth for her support and encouragement as well as my good friends Emma and Burkty for their friendship and time in assisting editing my papers.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF TABLES	v
LIST OF FIGURES	vi
ABSTRACT	viii
1 Introduction	1
2 Data Preparation	4
2.1 Data Collection and Preprocessing	4
2.2 Background Subtraction	6
2.2.1 Segmentation	6
2.2.2 Static Cluster Removal	10
3 Feature Extraction	14
3.1 Features	15
3.1.1 Number of Clusters	15
3.1.2 Center of Mass	17
3.1.3 Hu Moments	18
3.1.4 Extent and Solidity	18
3.1.5 Total Mass	21
3.1.6 Mass Hit	22
3.1.7 R & θ	24
3.1.8 Texture	26
3.2 Feature Selection	27

4	Machine Learning for Reaction Classification	29
4.1	Dataset	30
4.2	Experiments	30
4.2.1	Three Class Model with All Features	31
4.2.2	Two-Step Hierarchical Classification with All Features	31
4.2.3	Feature Removal	33
5	Deep Machine Learning for Reaction Classification	35
5.1	Pre-trained Remote Sensing Models	36
5.2	Transfer Learning with Fluorescence Microscopy Cell Imagery	38
5.2.1	Datasets and Preprocessing	40
5.2.2	Experimental Results	41
6	Autonomous SEM Imagery	45
6.1	Segment Anything	47
6.2	Mask Filtering	48
6.2.1	Texture	49
6.2.2	Experimental Results	50
7	Summary and concluding remarks	54
	BIBLIOGRAPHY	57
	VITA	63

LIST OF TABLES

Table	Page
3.1 Features chosen along with their comparison operator	27
4.1 Per-Class Model Performance with Averaged PFI	31
4.2 Hierarchical Classification Metrics	32
4.3 Classification Metrics with <i>Total Mass</i> and θ Features Removed	33
4.4 Step 2 Hierarchical Classification Metrics without <i>Total Mass</i> , θ , and R	34
5.1 Results of AI-NP segmentation through transfer learning on cell nuclei	41
5.2 ChangeFormer Based Reaction Classification Results	42
5.3 Hierarchical ChangeFormer Based Reaction Classification Results . . .	44
6.1 Texture Based Filtering Results	52

LIST OF FIGURES

Figure	Page
1.1 Reaction Examples: Spallation (a), Combination (b), and Sintering (c)	3
2.1 Autonomous Image Acquisition	5
2.2 Threshold Based Segmentation	7
2.3 Convex and Star-convex shapes	8
2.4 Example of Fluorescence Microscopy Cell Nuclei	9
2.5 Different Segmentation Methodology Results	9
2.6 StarDist Output Example	10
2.7 Example of Static and Changing Clusters	10
2.8 Difference Between the <i>before</i> and <i>after</i> Images of Figure 2.7	11
2.9 StarDist Based Statics Removal Model	12
2.10 Background Subtraction Process	13
2.11 Original Image Pair with its Background Subtracted Counterpart	13
3.1 Ideal Distance Matrix	15
3.2 Difference Matrix: Difference in Number of Clusters	16
3.3 Difference Matrix: Center of Mass Distance	17
3.4 Difference Matrix: Hu Moments Distance	18
3.5 Extent and Solidity Example	20

3.6	Distance Matrix: Extent(left) & Solidity(Right)	20
3.7	Distance Matrix: Total Mass	22
3.8	Mass Hit Example	23
3.9	Distance Matrix: Mass Hit	23
3.10	Feature Representation of R (a), when $i = 8$, and θ (b), when $k = 8$	24
3.11	Distance Matrix: R (left) & θ (Right)	25
3.12	Distance Matrix: R & θ	25
3.13	Distance Matrix: GLCM	26
3.14	Distance Matrix: All Features	28
3.15	t-SNE Embeddings of Proposed Features	28
4.1	Confusion Matrix for the Adaptive Boost model from Table 4.1	32
5.1	Dataset Samples: LEVIR(left) and DSIFN(right)	36
5.2	DSIFN Pre-trained Model Predictions	37
5.3	LEVIR RGB Versus Grayscale Input Predictions	38
5.4	ChangeFormer Modifications	39
5.5	Two different classes yielding identical change maps	43
6.1	Single Particle Reaction Image Modality Comparison	46
6.2	SEM Example Image	46
6.3	Segment Anything Prediction of Figure 6.2	48
6.4	Segment Anything Prediction with Grating Imperfections	49
6.5	LBP Image Example	50
6.6	Size Histogram for All Mask Regions	51
6.7	Histogram of All Exemplars: TPs(left) FPs(right)	52

ABSTRACT

Nanoenergetic materials offer high-density energy storage that may react to produce heat and release gaseous products. In order to utilize aluminum-based nanoenergetic materials effectively, the mechanisms by which aluminum fuel reacts to escape a passivating aluminum shell must be better understood. These mechanisms are affected by numerous parameters in both the laser-based photothermal heating setup to incite reactions as well as in the shape, size, and composition of the material itself. The vast number of parameters greatly increases the size of the search space when seeking a quantitative mapping between reaction parameters and reaction type. This thesis aims to accelerate this search through the creation of an autonomous image processing pipeline for nanoenergetic material reaction imagery, as well as a subsequent reaction classifier. The imagery is captured at 100x magnification utilizing optical microscopy and, various preprocessing methods are explored to aid in highlighting distinctive image features for characterizing reaction type. To perform the classification, we began with hand-crafted features and ensemble machine learning. Building on this, we modified a state-of-the-art transformer-based model for change detection to perform reaction classification. In doing so, we validated an abundant analog for nanoenergetic material, cell nuclei, to utilize for pre-training. Finally, we investigate texture-based image features of nanoenergetic material in scanning electron microscope (SEM) imagery. By leveraging this texture analysis, we can effectively isolate nanoenergetic material clusters within SEM imagery enabling autonomous laser targeting and imaging in an SEM environment.

Chapter 1

Introduction

Nanoenergetic materials have the potential to serve as a high-energy density storage mechanism. Certain nanoenergetic materials like aluminum nanoparticles (Al-NPs) can offer this at a low cost due to their general abundance. Al-NPs' low cost and high-energy density make them a good candidate as a component in nanothermite, a class of energetic composites consisting of heterogeneous solid-state fuel and oxide nanoparticles that are commonly applied in pyrotechnics, propellants, and explosives. Al-NPs are core-shell particles comprised of an aluminum fuel core surrounded by a passivating aluminum oxide shell with an average thickness of 2-4nm. Because of this, the aluminum core must escape the oxide shell to induce energy release. There are two main mechanisms that describe this escape: the diffusive oxidation mechanism (DOM) and the melt dispersion mechanism (MDM). Diffusion oxidation describes the thermally driven diffusion of aluminum and oxygen through the aluminum oxide shell with subsequent oxidation, which is slow and occurs at lower heating rates ($< 10^6 \frac{K}{s}$)[1][2]. This slower process typically causes a sintering reaction, leading to limited functionality and practicality in regards to the DOM. The melt dispersion

mechanism is much faster and transpires at higher heating rates ($10^6 - 10^9 \frac{K}{s}$) with volumetric expansion occurring as a result of this rapid heating causing a spallation of the aluminum oxide shell and rapid release of molten aluminum[3][4][5][6][7]. The MDM however, has not yet been significantly studied experimentally, and thus the heating parameters necessary to initiate it are not yet well understood. If the dependence on heating parameters is better understood, the application space of Al-NPs could expand and become more diverse.

Here, we focus on laser-based photothermally heated Al-NPs with different reaction types of spallation, sintering, or a combination of both. An example of each can be found in Figure 1.1. The spallation (Figure 1.1.a) is the more violent and higher energy reaction caused by the MDM. It typically causes the explosion of the cluster to expel molten material radially out from its center. The sinter (Figure 1.1.c) reaction is slower and less energetic. The material sinters together to form a smaller and denser cluster of particles. Finally, the combination reaction (Figure 1.1.b) is composed of both the spallation and sintering. From the example, it can be seen that some material is ejected away from the center cluster where the remaining material sinters together to form a more solid mass. The type of reaction is dependent on the physical and experimental parameters which include but are not limited to laser power density, pulse duration and wavelength, Al-NP shape and diameter, the quantity of particles within a cluster, and their proximity to one another. At present, there is a lack of quantitative mapping of these parameters to the corresponding reaction they elicit, and researchers manually characterize the reactions themselves through imagery captured before and after the reaction.

In this research, we focus on accelerating the search for the quantitative mapping between reaction parameters and reaction through creating a Al-NP reaction imagery pipeline or an autonomous reaction classifier. The pipeline can take in raw pre- and




Figure 1.1: Reaction Examples: Spallation (a), Combination (b), and Sintering (c)

post-reaction imagery of Al-NPs and process them into a state where an autonomous reaction classifier may label them. In the future, these two tools will greatly increase the speed at which Al-NP reaction data can be captured which is crucial to assemble a large dataset encompassing all physical and experimental reaction parameters. This grid of reaction labels and laser and cluster properties may then be used to search for a quantitative mapping between the properties and reaction labels.

Chapter 2

Data Preparation

2.1 Data Collection and Preprocessing

The imagery of spallation and sintering, or their combination, are organized as pairs of their pre- and post-reaction states. To create reactions, we utilized laser-based photothermal heating on Al-NPs deposited onto an optical grating substrate. The use of the grating allows for enhanced laser coupling with the nanoparticles resulting in the higher heating rates^[8] needed for spallation. After slide preparation, the imaging and laser irradiation are autonomous. A single image over the entire slide is captured via optical microscopy. This image is then processed utilizing ImageJ's included particle analysis feature^{[9][10]} which provides the coordinates for each Al-NP cluster on the slide. These coordinates can then be sent to our piezoelectric translation stage to align each cluster with the laser for irradiation. After all clusters have been irradiated, the stage then returns to the home position for the final post-reaction image. The process is shown in Figure 2.1. Finally, to extract an individual cluster's reaction image pair from the whole slide image, 380x380 pixel center crops were made around the targeted

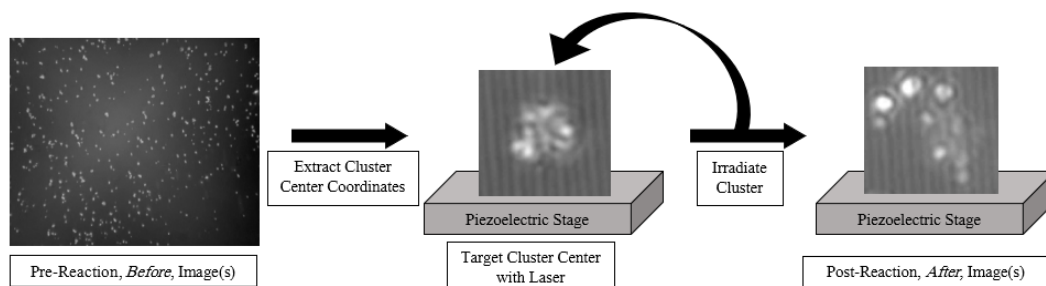


Figure 2.1: Autonomous Image Acquisition

clusters. We chose the patch size to be 30% larger than the maximum distance a mass of material was shown to be expelled from its cluster during spallation which was 145 pixels. This not only accommodated all imagery currently acquired but gave a sizeable buffer to allow for more violent reactions should they be encountered in the future. Ensuring all material was within the field-of-view (FOV) of the patch is crucial for accurate classification and mapping of laser parameters to reaction type.

The ordering of collecting only one image after all clusters had been irradiated caused a few problems. The first was that a particle could undergo a spallation, expelling material close to another cluster which caused inaccurate imagery for both clusters. This was fixed by removing the image reaction pairs for both clusters within 380 pixels of one another. Another issue was the piezoelectric's home position was not exact between the pre- and post-reaction image. To accurately register the images, an algorithmic registration utilizing intensity from Thévenaz et al.[11] was utilized. After filtering out all reaction pairs based off proximity, we were left with 159 reaction pairs: 61 Spallations, 73 Sinterings, and 25 Combination reactions.

2.2 Background Subtraction

The crucial aspect in accurately classifying a pair of pre- and post-reaction images into their reaction type lies in capturing and analyzing the differences or changes between the two images. To aid in highlighting the change between the two images, we sought to create a method of background subtraction. The background subtraction would not only remove the optical grating in the background but also any cluster within the image FOV not involved in the reaction. The background subtraction method had to be devised within a low-data environment and was separated into two parts: segmentation and static cluster removal. The segmentation would be used to isolate all clusters, statics, and non-statics from the background grating while the cluster removal determined which clusters were not involved within the reaction and therefore removed.

2.2.1 Segmentation

The first methodology for segmentation pursued was simple thresholding and grayscale morphology. Because our imagery features a light target on a dark background, we saw some success from this approach. To begin, we applied contrast limited adaptive histogram equalization (CLAHE) [12] to both the before and after reaction image. This was done to increase the contrast in the image and aid in equalizing the dark spots within a cluster. Afterward, a Gaussian blur was applied to smooth out rougher corners and edges. We then calculated a histogram over each image and chose the thresholding value to be a constant 60-pixel value (0-255) higher than the most populated bin. This value was chosen experimentally and optimized to remove more content as erroneous false positive material left behind in an image causes more harm than lost material. Finally, a morphological closing was applied to aid in closing




Figure 2.2: Threshold Based Segmentation

gaps within the clusters and erroneous small regions left in bright parts of the optical grating.

An example of the final result compared with the original imagery can be seen in Figure 2.2. The figure shows some of the fringing in the bottom cluster of the *After* image was left behind in the segmented image. While this imperfection and imperfections like it in other images are not ideal, the real downfall of this technique is its inconsistency from image to image. Many images had significant cluster regions of interest entirely removed while others had large areas of background segmented. Our dataset had significant brightness fluctuations from image to image and while CLAHE helped mitigate this, it was clear an intensity-based segmentation process was not ideal.

To improve our methodology for segmentation, we discovered a unique segmen-




Figure 2.3: Convex and Star-convex shapes

tation model called StarDist[13]. While being a deep convolutional neural network (CNN), StarDist utilizes a lightweight U-Net[14] architecture reducing the total number of trainable parameters allowing effective training on smaller datasets. StarDist performs non-maximum suppression with star-convex polygons instead of traditional bounding boxes to match or outperform larger models like Mask R-CNN[15] in crowded contexts given that the segmentation target is star-convex. For a shape to be convex, any two points within the object must be connectable by a straight line without leaving the object’s area. Star-convex is similar, however, there must be at least one point within the object that can reach any other without leaving (Figure 2.3). Fortunately, the vast majority of the clusters in our imagery are star-convex.

To train the StarDist model for nanoenergetic material segmentation, we hand-labeled each pre- and post-reaction image creating binary masks as our ground truth. The pre- and post-reaction imagery was combined into the same group as there are only slight differences in appearance between reacted and non-reacted Al-NPs in optical imagery. StarDists supplies a model pre-trained on fluorescence microscopy imagery of cell nuclei which we utilized. Fluorescence microscopy imagery of cell nuclei also contains a light target on a dark background making the pre-trained StarDist model well suited for transfer learning to Al-NP imagery. We refined the pre-trained




Figure 2.4: Example of Fluorescence Microscopy Cell Nuclei

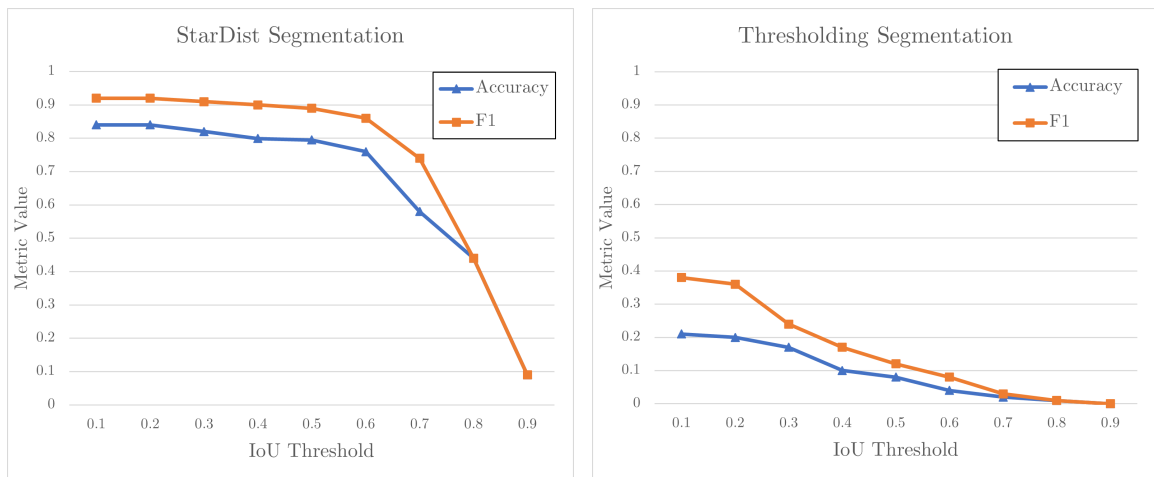


Figure 2.5: Different Segmentation Methodology Results

StarDist for 100 epochs on 90% of our imagery. Figure 2.5 displays the results of this refined model on the remaining 10% of our dataset, which was not used for training, across various IoU thresholds. Not only does the StarDist-based segmentation increase the overall metric performance in seeing clusters and removing background, but it also has good, steady performance up until 0.6 IoU leading to much better contours.




Figure 2.6: StarDist Output Example




Figure 2.7: Example of Static and Changing Clusters

2.2.2 Static Cluster Removal

Now that we are able to separate all clusters from the background, we needed a process that could remove clusters not involved in the reaction. Clusters like this should be static and unchanged from pre- to post-reaction state. Figure 2.7 shows static clusters circled in red and changing clusters boxed in green. Static clusters clutter many of the reaction pairs in our dataset and would heavily impact any visual features extracted from the imagery. Because of the success we had using StarDist for segmentation, we chose to use it again for static cluster removal.

One option was to simply use the *before* and *after* images of a reaction pair as two different channels for input into the model. However, because we already had an accurate method of registration and we know that static blobs are mostly




Figure 2.8: Difference Between the *before* and *after* Images of Figure 2.7

unchanging, we chose to take the difference between the *before* and *after* images within each reaction pair. This difference should highlight parts of the image that are changing, while unchanging regions should trend towards 0. Figure 2.8 shows the difference between the *before* and *after* images of Figure 2.7 as well as the mask for only the changed area. The areas of deeper color signify heavily altered regions whereas unchanging regions are close to white. Because of this success, we hand-labeled them for ground truth masks and supplied the difference images to a new, randomly initialized StarDist model for training. Like before, we trained with 90% of our dataset but for slightly longer, 150 epochs, as we did not use the pre-trained model. After training, we tested the model on the remaining 10% of the dataset and once again saw good results as shown in Figure 2.9.

With a model that can accurately segment AI-NPs and another which can detect static clusters, we can now combine them to form the final background subtraction process. Referencing Figure 2.10, both the original pre- and post-reaction images (a) are first processed through the segmenting model yielding two cluster masks (b) and secondly, subtracted from one another resulting in the difference image for cluster filtering(c). Next, the difference image is run through the static cluster removal model which outputs a mask of only clusters that are involved in the reaction (d).

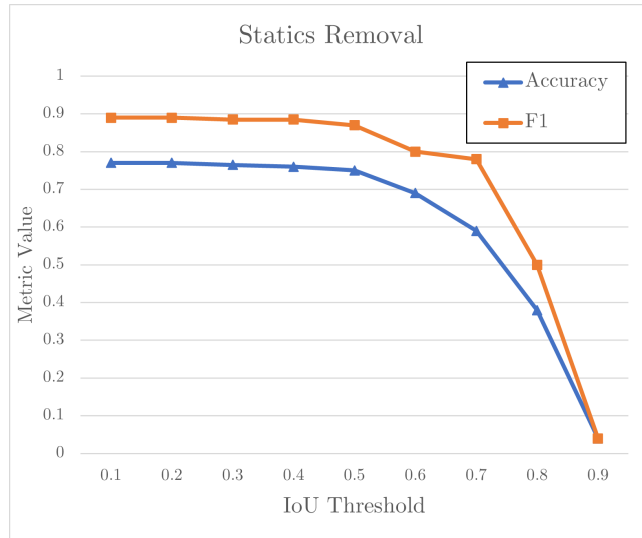


Figure 2.9: StarDist Based Statics Removal Model

Finally, the non-static cluster mask can be pixel-wise logical OR'd with both the *before* and *after* cluster masks which generates our final two masks (e). These final two masks can be applied to the original imagery to receive the final image pair with the background removed, Figure 2.11. Ideal for feature extraction, the background subtracted imagery no longer contains any erroneous data, preserving the integrity of underlying patterns within the extracted features.

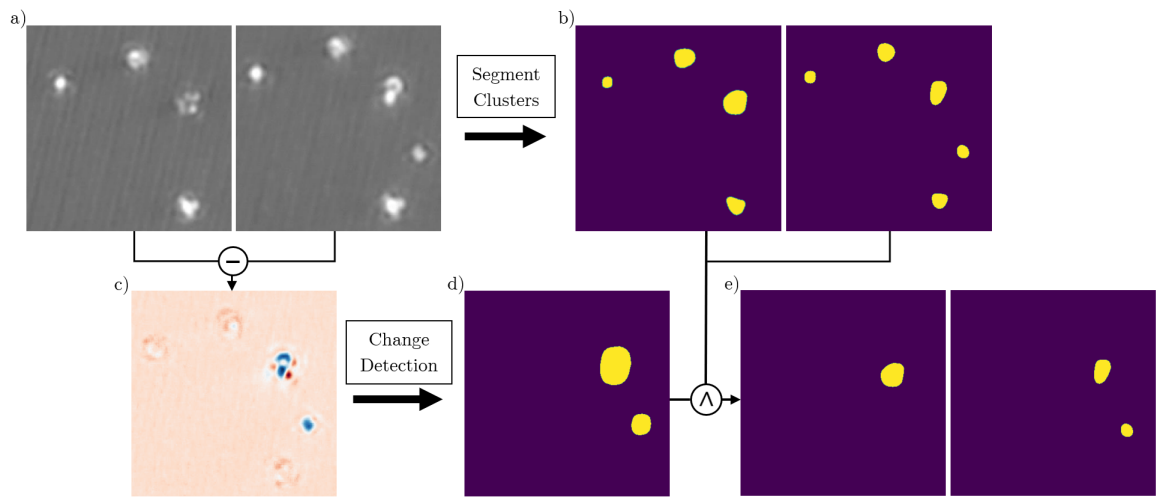


Figure 2.10: Background Subtraction Process




Figure 2.11: Original Image Pair with its Background Subtracted Counterpart

Chapter 3

Feature Extraction

With a preprocessed dataset, we can now focus on feature extraction. The aim of this chapter is to discuss distinctive image features that may be used to classify a reaction image pair into one of the following reaction classes: sintering, spallation, or combination. A feature extractor will be applied to both the *before* and *after* images and compared using a comparison operator. The comparison operator will depend on the feature extracted but could be a subtraction, ratio, or distance metric. The compared feature, which is now a scalar, is added to the reaction pair's feature vector. The length of the feature vector is dependent on the number of feature types extracted from the image. When examining a single feature, the feature vector has a length of 1, examining two features will give a length of 2, and so on. To evaluate a feature or list of features, we utilize a distance matrix. The ideal distance matrix is shown in Figure 3.1. The distance matrix is a heatmap showing the distances between all reaction pairs' feature vectors. For example, in Figure 3.1, the highlighted cell illustrates the distance between the feature vector of the second sintering reaction pair and that of the third combination reaction pair. Effective features, or feature sets, will exhibit

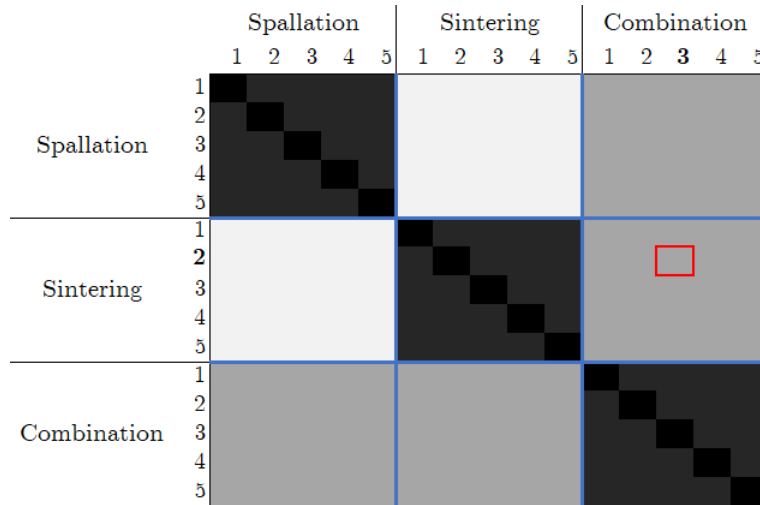


Figure 3.1: Ideal Distance Matrix

dark regions in areas of intra-class comparison. These dark areas signify that the distances between feature vectors of the same class are small and alike. In contrast, for inter-class comparison, like between sintering and spallation, we would like to see a lighter-colored area showing that feature vectors from different classes have a greater distance from one another. The combination class is special. For features or feature sets to be effective here, we need to see dark regions for intra-class comparison but only somewhat light-ish areas for inter-class comparison. Combination reactions exhibit characteristics from both spallations and sinterings so it makes sense if their distances in inter-class scenarios are a bit lower.

3.1 Features

3.1.1 Number of Clusters

The first feature examined is the difference of the number of clusters between the pre- and post-reaction images. Spallations, which are the more violent reaction, often




Figure 3.2: Difference Matrix: Difference in Number of Clusters

eject molten material from the cluster center which should increase the total number of clusters after reaction. One possible issue with this feature is when the reaction is so violent that the molten material ejected is too small to be resolved by the imaging sensor. This leaves the *after* image completely blank. The phenomenon is quite common in our experimental set-up but is easily remedied as it is only possible in spallations or in rare cases of severe preprocessing inaccuracies. If a blank *after* image is encountered, then the reaction pair’s value is set to the 90th percentile of all values in this feature as we want the value to be large, but not set to an outlier. From the difference matrix in Figure 3.2, we see decent results for inter- and intra-class comparison for sinters and spallations. A higher inter-class result for combinations and spallations is not an issue as it furthers the separation between the two classes. It also signifies that for this feature, combinations are more similar to sinters than spallations. And, therefore, combinations have fewer ejections on average than spallations.

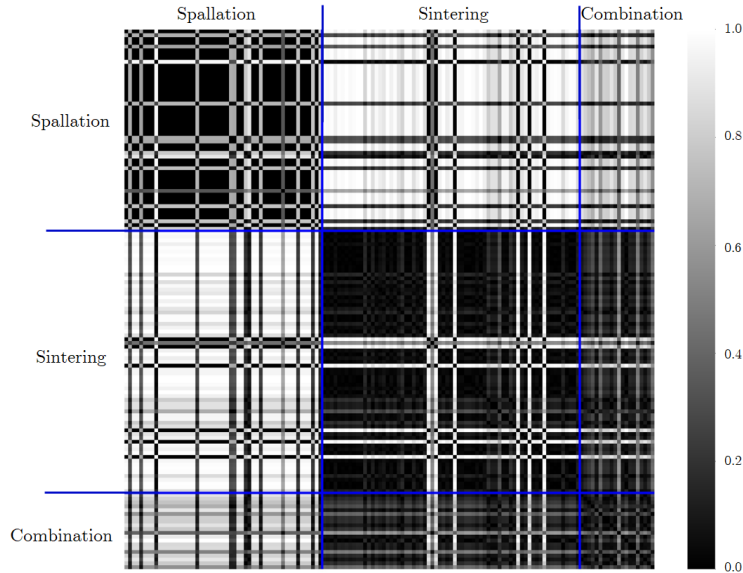


Figure 3.3: Difference Matrix: Center of Mass Distance

3.1.2 Center of Mass

To capture the center of mass from an image, we simply abstract mass to be the pixel intensity values. We then find the center of the pre- and post-reaction images weighted by the pixels' intensities and find the Euclidean distance between the two points. Because spallations are more violent, we expect the highest distances to be from spallations, then the combination class, and finally the sinters. Like before, if the *after* image is blank, the distance is set to the 90th percentile. Figure 3.3 shows excellent separation for the spallation class. Again, we see the sintering and combination class to be more similar but at least still slightly brighter than the intra-class sintering comparisons. Interestingly, the intra-class combination comparisons do not appear to be similar to one another. While a single feature yielding sub-optimal results for the combination class is not a major concern, if multiple features that exhibited similar behavior were utilized, it could cause the combination class to become inseparable from other classes.




Figure 3.4: Difference Matrix: Hu Moments Distance

3.1.3 Hu Moments

In theory, the shape of the clusters in the combination reactions change much more than those in the sintering class. Hu moments[16] are a translation, scale, rotation, and reflection invariant shape descriptor. By computing the Hu moments over both the *before* and *after* images and comparing the vector outputs utilizing Euclidean distance, we can characterize how much the shape changed from the reaction. When analyzing the difference matrix for Hu moments, Figure 3.4, all classes seem to blend together except for the combination class. The entries in the combination class have little likeness in inter- or intra class comparisons. Because of these reasons, Hu moments are not helpful for characterizing reaction pairs.

3.1.4 Extent and Solidity

While characterizing shape changes in reaction outcomes is straightforward, when dealing with singular regions of interest, it becomes more challenging when multiple

regions of interest are present in an image. This situation is common in the *after* images within our dataset. A better descriptor related to shape is the distribution of the clusters and how it changes from the reaction. Two common measurements for this are extent and solidity. Extent is the ratio of target pixels to the area of the target's bounding box. Similarly, solidity is the ratio of target pixels to the area of the target's convex hull. Figure 3.5 shows a nice visual representation of these two features. Because spallations and combinations both eject material from the cluster center, the extent and solidity should drop significantly from these reactions. Conversely, sintering reactions should have little changes to extent and solidity between pre- and post-reaction. We chose subtraction as our comparison operator for both features so that high values correlate with significant drops in measurements whereas low values denote little change. Like we had done in other features, if the *after* image contains no material, the feature value is set to the 90th percentile. Both features offer great spallation versus sintering separation and sintering intra-class likeness. The sintering and combination intra-class similarity is slightly higher for solidity, making it the preferred choice.




Figure 3.5: Extent and Solidity Example

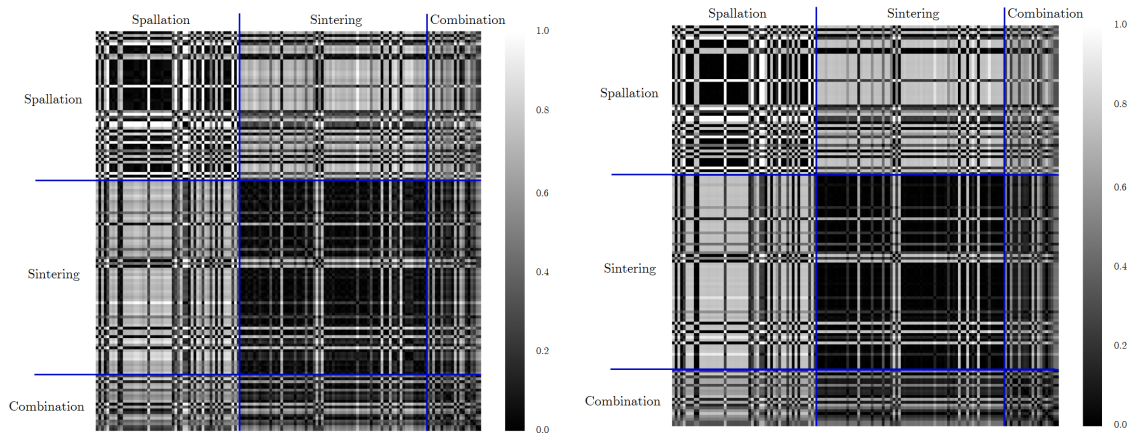


Figure 3.6: Distance Matrix: Extent(left) & Solidity(Right)

3.1.5 Total Mass

Continuing on with pixel intensity as an abstraction for cluster mass, we chose to extract the total mass within the image pair. The philosophy came from some visual observation that combination reactions generally take place in larger clusters, at least for the current laser properties. Additionally, spallations are associated with losing mass due to some ejections being too small for our imaging sensor to resolve. In contrast, sintering reactions should have a constant or increased mass due to them typically losing little to no area and an increase in intensity. In contrast, sintering reactions are expected to exhibit constant mass or even small gains, as they typically experience minimal area loss and an increase in intensity. We utilized the ratio of the measurements from the two different time domains as our comparison operator for this feature. Total mass shows great intra-class likeness and inter-class separability for spallations. Again, we see some overlap for sinters and combinations which, while not helpful, would not hurt overall performance as long as there are some features that do show separability in those classes.




Figure 3.7: Distance Matrix: Total Mass

3.1.6 Mass Hit

Extrapolating on the idea of total mass, instead of examining all the mass in the image, we can focus only on the area the laser hits. This allows the feature to capture both the material loss or gain and the positioning of clusters between the *before* and *after* images. An example of this feature is seen in Figure 3.8. In this figure, the white ellipse represents the laser's beam, and only the mass within the area is summed for this feature. Like total mass, we again used the ratio of the two measurements for the comparison operator. Mass hit's distance matrix, shown in Figure 3.9, is quite interesting. It shows excellent results for intra- and inter-class comparisons for spallations, but seemingly useless results for sinters and combinations.




Figure 3.8: Mass Hit Example




Figure 3.9: Distance Matrix: Mass Hit

3.1.7 R & θ

Both features R and θ are extracted from the Fourier Transform of the imagery. These two features, introduced by Li et al[17], were shown to be effective for palmprint identification. The R feature is the i individual sums of pixel responses within expanding rings in the transformed image, Figure 3.10a. Similarly, the θ feature is the sums within the k fanning wedge shapes, Figure 3.10b. While in palmprint identification, these features are used to extract crease strength and orientation. The features' main purpose here is to illustrate the strength of a reaction. This is captured by taking the Euclidean distance between the measurements received from the *before* and *after* images. Figure 3.11 shows the distance matrices for both R and θ . Although R and θ 's distance matrices demonstrate many similarities, a few key differences are present. In particular, the intra-class likeness and inter-class separability for combinations are better for the θ feature. On the other hand, the intra-class likeness for the sinters seems to be better in the R feature. The result of combining the two features together is shown in Figure 3.12, which offers the characteristics of both the previous distance matrices.




Figure 3.10: Feature Representation of R (a), when $i = 8$, and θ (b), when $k = 8$

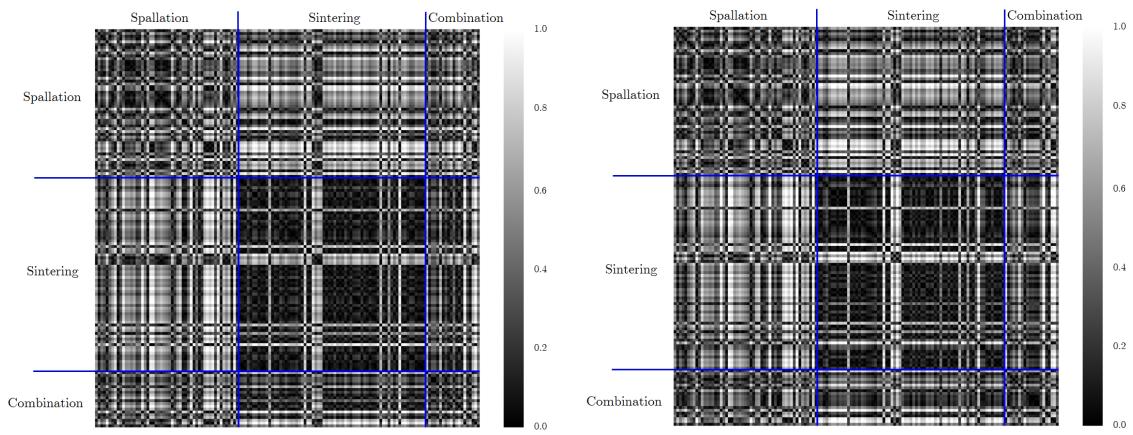


Figure 3.11: Distance Matrix: R (left) & θ (Right)




Figure 3.12: Distance Matrix: R & θ

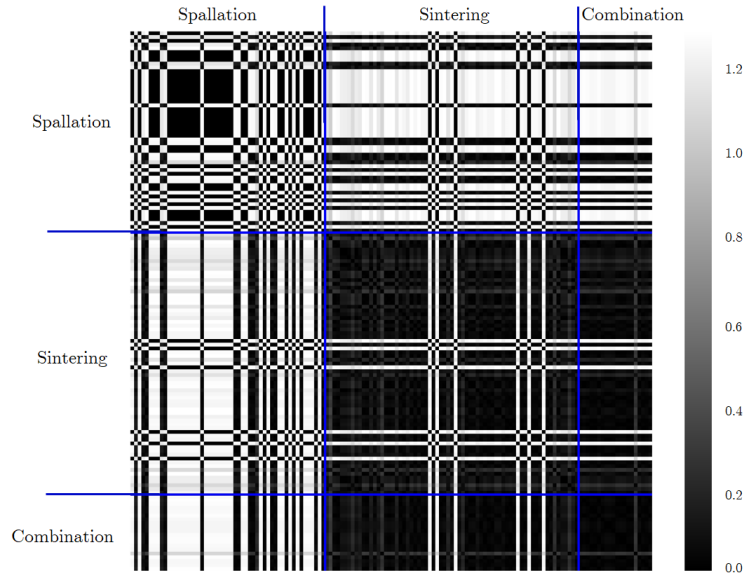


Figure 3.13: Distance Matrix: GLCM

3.1.8 Texture

The last feature examined was texture. We were examining whether the different reaction types affected the post-reaction cluster texture differently. To do this, we extracted the gray level co-occurrence matrix[18] (GLCM) with varying distances and eight equally spaced angles over the clusters in the *after* images only. Next, we extracted the five texture properties of contrast, dissimilarity, homogeneity, energy, and correlation from the co-occurrence matrices. These five texture properties were used as the feature values. Again, no comparison operator is needed for this feature as we are not concerned with how the texture changed but more of what the texture changed into. Unfortunately, this texture measure is not effective on our data as seen in Figure 3.13. While this is only one of many possible texture descriptors to examine, we chose to move forward as the optical imagery clearly does not have spatial resolution to resolve minute texture differences. Future work involving imagery from scanning electron microscopes (SEMs) may revisit texture as a viable descriptor.

3.2 Feature Selection

After analyzing the features, we chose to move forward with seven of them. The features chosen along with their comparison operator can be seen in Table 3.1. To

Feature	Comparison Operator
<i>Center of Mass</i>	Euclidean Distance
<i>Number of Clusters</i>	Subtraction
<i>Solidity</i>	Subtraction
<i>Mass Hit</i>	Ratio
<i>Total Mass</i>	Ratio
<i>R</i>	Euclidean Distance
<i>θ</i>	Euclidean Distance

Table 3.1: Features chosen along with their comparison operator

further analyze how all the features work together, we concatenated all features into a single feature vector for each reaction pair. The distance matrix for all the features together can be seen in Figure 3.14. From this distance matrix, we can see excellent inter- and intra-class performance for spallations, a common theme during feature exploration. The inter-class performance between combinations and sinters is not nearly as good but their intra-class performance is fair.

Another way to examine our features is with T-distributed Stochastic Neighbor Embedding[19] (t-SNE). Two different sample embeddings that show possible separability based on our proposed features can be seen in Figure 3.15. The embeddings echo the probable class overlap between the combinations and the sinters. These seven features may now be used to create machine learning models for reaction classification.

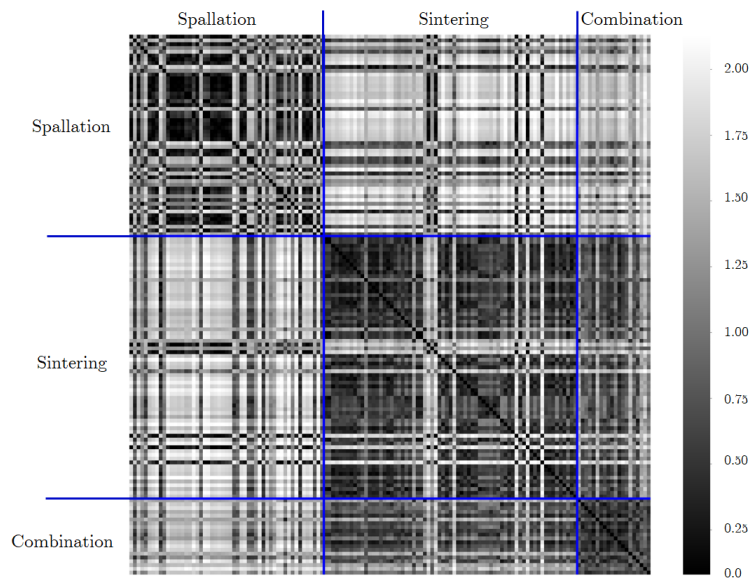


Figure 3.14: Distance Matrix: All Features

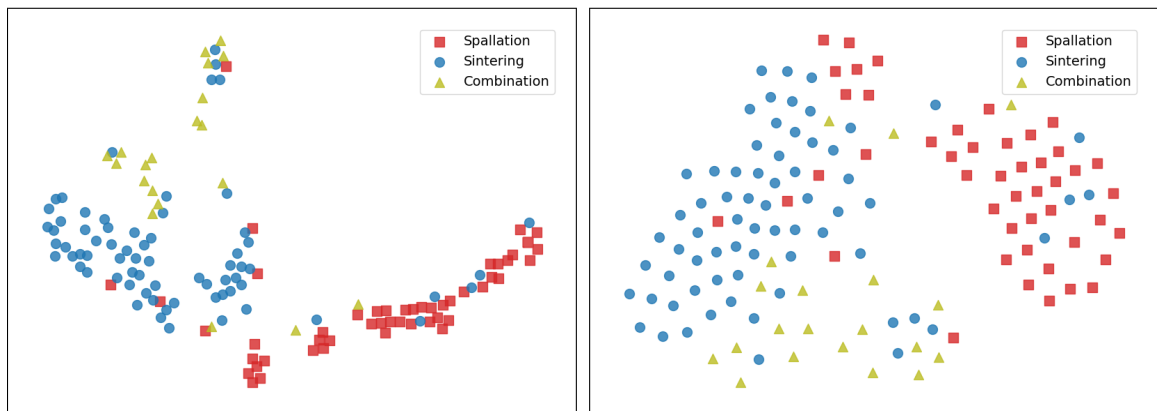


Figure 3.15: t-SNE Embeddings of Proposed Features

Chapter 4

Machine Learning for Reaction Classification

Experiments are performed to create the best reaction classifier using the feature vector created in Section 3.2. To accomplish this, we performed hyperparameter grid searches over the following classifiers: Random Forest[20], Adaptive Boosting[21][22], Gradient Boosting[23], and Bagging[24]. Each model was built using Scikit-learn[25], and the ensemble models use a decision tree as their base classifier. In addition to analyzing model performance, we conduct feature permutation tests to identify the most important features and to exclude features limiting results. We chose to begin with ensemble models and only features as input instead of deep neural networks and full image input. This decision was prompted by our small dataset and to serve as a foundational baseline to examine what performance could be gained with more complex implementations.

4.1 Dataset

All models were trained with a custom dataset utilizing imagery captured through the process outlined in Chapter 2. Using this process, we collected 159 reaction pairs: 61 Spallations, 73 Sinterings, and 25 Combination reactions. To reinforce our relatively small dataset, we generated additional feature vector samples by augmenting existing reaction pairs. The augmentation strategies used 90° rotations of the images, as well as horizontal and vertical image flips that result in increasing the training samples by $12\times$. Finally, to increase model robustness to image brightness, we incorporated a uniformly random brightness augmentation of $\pm 25\%$ to 65% of training samples.

4.2 Experiments

To evaluate all of our chosen models, we ran our hyperparameter grid search for each model type using 5-fold cross-validation. During the hyperparameter search, each combination of hyperparameters was trained on four different training folds while leaving one fold out for validation. The process was repeated 5 times while cycling which fold was used for validation. The validation scores were then averaged to give a final scoring metric for the combination of hyperparameters. For each model type, the best five hyperparameter combinations were then retrained on all five training folds and evaluated on a sequestered test set held out of training in all previous steps. Additionally, the test set was used for scoring feature importance. The feature importance measurements across all hyperparameter combinations within each model type are then averaged.

Model	Test Set Performance			Test Set Averaged Permutation Feature Importance							
	F1	F1	F1	<i>Center of Mass</i>	<i>Number of Clusters</i>	<i>Solidity</i>	<i>Mass Hit</i>	<i>Total Mass</i>	<i>R</i>	θ	
	Sinter	Combination	Spallation								
Random Forest	0.753	0.393	0.902	[HTML]FCFF2F 0.297	-0.002	0.092	-0.021	-0.090	-0.022	-0.049	
Adaptive Boost	0.869	0.591	0.957	[HTML]FCFF2F 0.273	0.069	0.117	0.084	-0.005	0.062	-0.011	
Gradient Boost	0.850	0.515	0.949	[HTML]FCFF2F 0.327	0.032	0.160	0.049	-0.026	0.009	-0.004	
Bagging	0.845	0.524	0.957	[HTML]FCFF2F 0.373	0.022	0.093	0.011	-0.012	-0.001	0.000	

Table 4.1: Per-Class Model Performance with Averaged PFI
1st, **2nd**, **3rd** Most Important Features

4.2.1 Three Class Model with All Features

Table 4.1 shows the performance metrics on the test set for the models trained on all folds. Because permutation feature importance[20] is defined as a decrease in the model’s scoring when the feature is randomly shuffled, any cells showing a negative score signify that the random shuffling increased the model’s score. These negative scores may indicate issues such as overfitting or that feature is irrelevant or even problematic for the classification task. Consistently seeing the same feature showing low or negative scores across multiple model types like *Total Mass* and θ in Table 4.1 strongly suggests they may be problematic. Another issue is the low F1 scores for the Sintering and Combination class. Looking back at the t-SNE embeddings from Figure 3.15, we can visually see the Sintering and Combination classes exhibit more overlap compared to the Spallations. Examining the confusion matrix shown in Figure 4.1 further shows the extent to which these two classes overlap.

4.2.2 Two-Step Hierarchical Classification with All Features

We then moved to a two-step hierarchical classification approach to address this significant overlap. The first step separates the Spallations from the Sinters and Combination reactions, whereas step two separates the Sinters from the Combinations. Results for this implementation can be seen in Table 4.2. From these results, it can be seen that the classification task of Spallation vs Other (Sinters and Combinations)

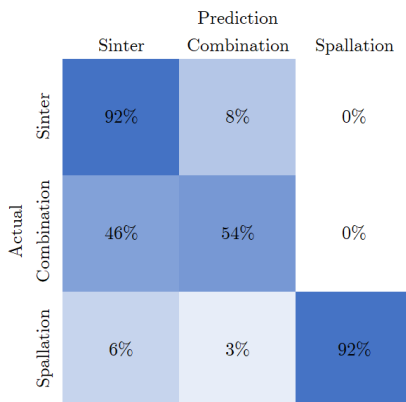


Figure 4.1: Confusion Matrix for the Adaptive Boost model from Table 4.1

Model	Test Set Performance			Test Set Averaged Permutation Feature Importance						
	F1 Sinter	F1 Combination	F1 Spallation	<i>Center of Mass</i>	<i>Number of Clusters</i>	<i>Solidity</i>	<i>Mass Hit</i>	<i>Total Mass</i>	<i>R</i>	<i>θ</i>
Random Forest	0.926	0.953	[HTML]9B9B9B	-0.034	0.007	[HTML]FFFF00 0.038	-0.007	-0.046	0.004	-0.081
Adaptive Boost	0.979	0.987	[HTML]9B9B9B	0.026	0.014	[HTML]FFFF00 0.060	0.028	-0.020	0.002	-0.009
Gradient Boost	0.953	0.972	[HTML]9B9B9B	0.016	0.017	[HTML]FFFF00 0.076	0.033	-0.039	0.003	-0.027
Bagging	0.966	0.980	[HTML]9B9B9B	0.001	0.011	[HTML]FFFF00 0.047	0.015	-0.021	-0.002	-0.040

Table 4.2: Hierarchical Classification Metrics
 Spallation vs Other (top) and Sinter vs Combination (bottom)
1st, **2nd**, **3rd** Most Important Features

is much simpler than the three class task as expected. From a feature importance perspective, the *Total Mass* and θ features are no longer mostly negative and the θ feature is even the third most important for the Bagging models. When looking at the F1 scores for Sinter vs Combination models, we see either small losses or large leaps forward when compared to the three class models of Table 4.1. Another benefit to separating the classification tasks is that we can see further specifications of what features are important for which reactions. For example, the Sinter vs Combination models are the first where we did not see high importance for the *Center of Mass* feature which shows that it is primarily useful for determining Spallations. Lastly, we once again see many negative values for the *Total Mass* and θ features and chose to omit them for the next round of training.

4.2.3 Feature Removal

Results of models trained with suspected problematic features removed are shown in Table 4.3. The F1 score for Spallation in the three class models decreased for every model type except for the Random Forest. Even though there are small decreases in performance, ≤ 0.012 , because the drop is prevalent in most of the models, it is likely that the removed features *Total Mass* and θ are useful for identifying Spallations. This hypothesis is further supported by the positive feature importances in the top half of Table 4.2. F1 scores for Sinterings and Combinations in the three class models show predominantly increases in performance which supports the hypothesis that the removed features were hindering performance for those classes. The hierarchical approach exhibits highly significant performance increases in both classes' F1 scores, affirming it as the superior classification methodology for this task. In both classification methodologies, we see increased negative feature importance measurements for the *R* feature.

The last experiment with ensemble machine learning models tests the efficacy of the *R* feature for the classification of Sinterings and Combinations. Like before, we retrain models without *Total Mass* and θ , and now remove the *R* feature as well. Table 4.4 shows this experiment improved the Adaptive and Gradient Boosting models' score but decreased the Random Forest and Bagging models. This observation

Model	Test Set Performance			Test Set Averaged Permutation Feature Importance				
	F1	F1	F1	<i>Center of Mass</i>	<i>Number of Clusters</i>	<i>Solidity</i>	<i>Mass Hit</i>	<i>R</i>
	Sinter	Combination	Spallation					
Random Forest	0.829	0.514	0.910	[HTML]FFFF00 0.273	0.055	0.146	-0.035	-0.012
Adaptive Boost	0.878	0.587	0.945	[HTML]FFFF00 0.266	0.084	0.144	0.058	0.022
Gradient Boost	0.890	0.596	0.941	[HTML]FFFF00 0.283	0.045	0.188	0.085	-0.003
Bagging	0.907	0.675	0.952	[HTML]FFFF00 0.368	0.013	0.142	-0.003	-0.002
Random Forest	0.899	0.701	[HTML]9B9B9B	0.026	0.060	[HTML]FFFF00 0.078	0.026	0.023
Adaptive Boost	0.929	0.764	[HTML]9B9B9B	0.063	0.033	[HTML]FFFF00 0.095	0.018	-0.005
Gradient Boost	0.936	0.782	[HTML]9B9B9B	0.025	0.017	[HTML]FFFF00 0.108	0.015	-0.001
Bagging	0.932	0.727		0.009	[HTML]FFFF00 0.043	0.047	0.038	0.003

Table 4.3: Classification Metrics with *Total Mass* and θ Features Removed
 Three Class Model (top) and Step 2 of Hierarchical Approach (bottom)
 1st, 2nd Most Important Features

Model	Test Set Performance		Test Set Averaged Permutation Feature Importance			
	F1	F1	<i>Center of Mass</i>	<i>Number of Clusters</i>	<i>Solidity</i>	<i>Mass Hit</i>
	Sinter	Combination				
Random Forest	0.822	0.579	0.001	0.002	[HTML]FFFF00 0.067	0.008
Adaptive Boost	0.944	0.790	[HTML]FFFF00 0.073	0.033	0.093	0.026
Gradient Boost	0.943	0.800	0.041	0.010	[HTML]FFFF00 0.136	0.014
Bagging	0.928	0.725	0.019	[HTML]FFFF00 0.034	0.045	0.025

Table 4.4: Step 2 Hierarchical Classification Metrics without *Total Mass*, θ , and *R* **1st**, **2nd** Most Important Features

matches with the feature importance measurements seen in Table 4.3, as Adaptive and Gradient Boosting are the models that had negative values for the *R* feature. Judging by the fact that the Random Forest and Bagging classifiers have greater depth in their estimators, one possible explanation for this observation is that the *R* feature requires this depth in order for it to be useful. Regardless, the two best models for Sintering and Combination classification of all experimentation thus far are found in this last experiment. These models will serve as a performance baseline to compare against future, more complex models. Additionally, these models help verify our proposed hand-crafted features are useful in the characterization of reaction imagery.

Chapter 5

Deep Machine Learning for Reaction Classification

Previously, we relied on creating handcrafted features extracted from our imagery coupled with ensemble machine learning models like Boosting[21], [22], [23] and Bagging[24] to classify the reactions. Here, we focus on adapting a state-of-the-art transformer-based siamese network for change detection, ChangeFormer[26], to perform reaction classification. Through the use of ChangeFormer, we aim to utilize not only our proven handcrafted features but also leverage the raw imagery itself.

ChangeFormer is composed of a shared-weight siamese encoder, difference modules, and a decoder. Its encoder performs multi-scale feature extraction at four different resolutions on both the *before* and *after* images. Each of the four pairs of multi-scale features is then run through their difference module before being concatenated and fused together for use in the decoder to generate a binary change map.

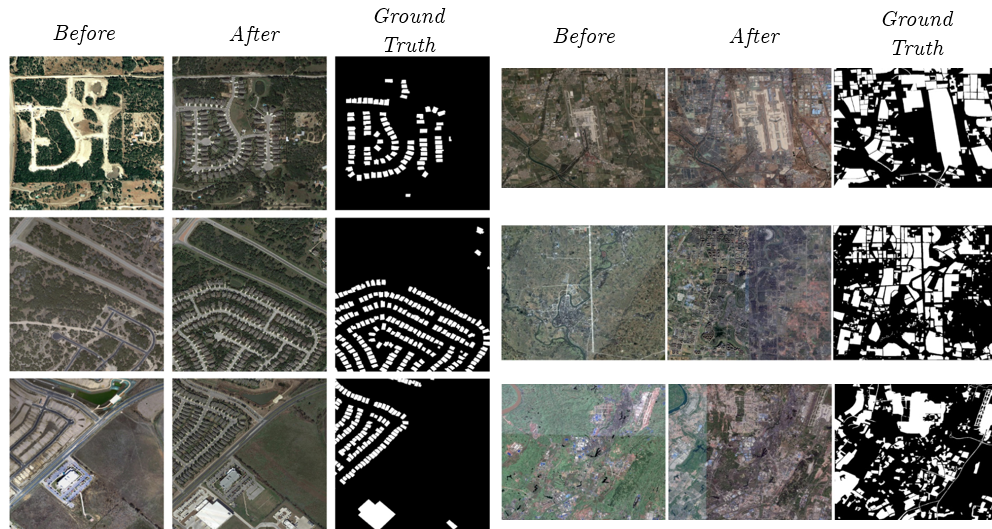


Figure 5.1: Dataset Samples: LEVIR(left) and DSIFN(right)

5.1 Pre-trained Remote Sensing Models

The authors of ChangeFormer have released two different pre-trained ChangeFormer models. One model was trained on LEVIR-CD[27], a change detection dataset primarily focused on residential and commercial buildings. The other model was trained on DSIFN[28], a high-resolution dataset captured at much higher ground sampling distance. This dataset focuses on changes in land cover objects. Example imagery from both datasets is shown in Figure 5.1.

To begin, we decided to run some reaction pairs through each pre-trained model. Because these pre-trained models require RGB imagery and our dataset is grayscale, we simply duplicated our one-band twice to create three-banded imagery. While the pre-trained LEVIR model had completely blank predictions on our imagery, the DSIFN model had some interesting results, Figure 5.2. Because both the LEVIR and DSIFN datasets are so different from our imagery, it’s not surprising the models do not predict accurate change maps for our imagery. What is strange, however, is how the predictions from the DSIFN model shown in Figure 5.2 appear close to an




Figure 5.2: DSIFN Pre-trained Model Predictions

accurate prediction, but inverted. As a quick experiment, we took imagery from the original LEVIR dataset and transformed it into a single-banded grayscale. We then duplicated that single grayscale band like we did with our AI-NP imagery to form a three-banded grayscale. Figure 5.3 shows how grayscale three-banded imagery destroys the model’s prediction, showing that the model heavily relies upon ‘real’ RGB three-banded imagery. Because of this, we chose to train our ChangeFormer model from scratch rather than refine the pre-trained models.

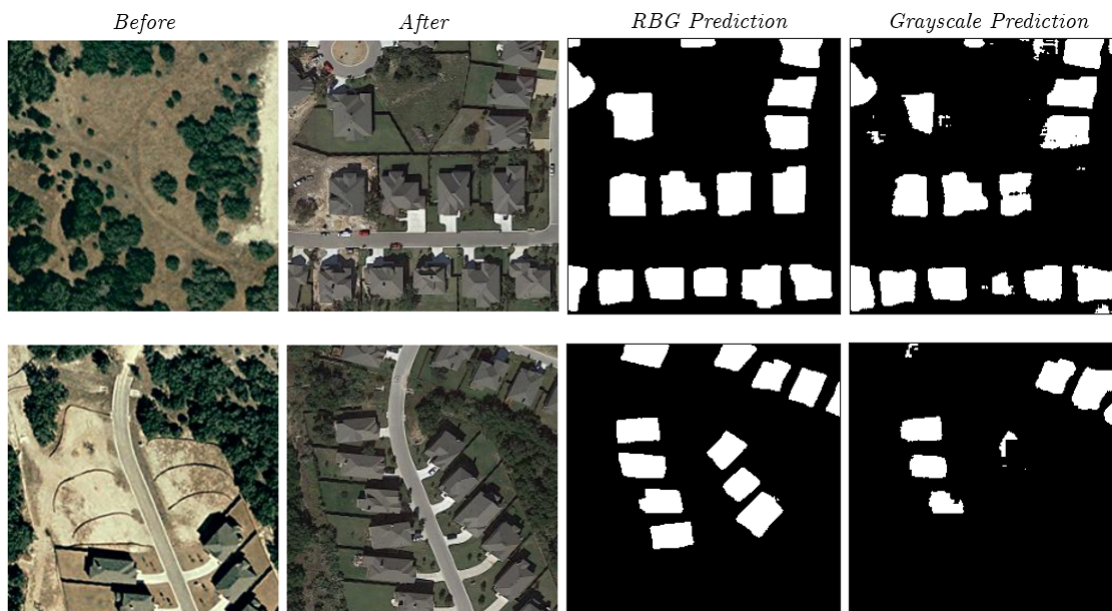


Figure 5.3: LEVIR RGB Versus Grayscale Input Predictions

5.2 Transfer Learning with Fluorescence Microscopy Cell Imagery

Because of ChangeFormer’s large size and our smaller dataset, we first decided to pre-train ChangeFormer’s encoder to segment fluorescence microscopy images of cell nuclei. Working with cell nuclei for pre-training has been shown to transfer well to our AI-NP imagery in chapter 2. After the encoder was trained on cell imagery, we refined the encoder using our AI-NP images and can transfer this trained encoder into a full ChangeFormer model. Furthermore, this encoder can serve as an enhanced AI-NP segmentor, offering improved accuracy in regenerating our handcrafted features. As an optional secondary pre-training step, we can use this ChangeFormer model with the refined encoder to generate binary change maps. This will allow components deeper into the model, an example being the difference modules, to be pre-trained using a task adjacent to reaction classification.

To alter ChangeFormer from predicting a change map to a reaction classification of

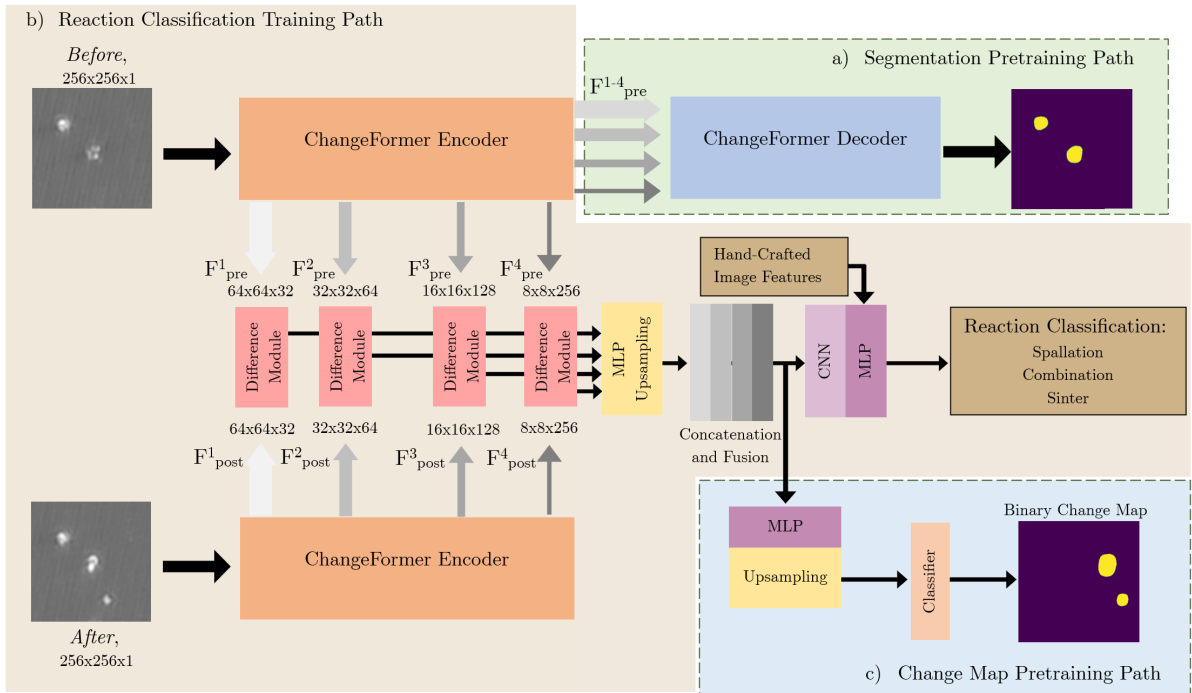


Figure 5.4: ChangeFormer Modifications
a) Segmentation Pre-training, b) Reaction Classification, c) Change Map Pre-training

Spallation, Sintering, or Combination, we removed the upsampling after the concatenation and fusion and inserted a CNN with a multi-layer perception (MLP) classifier. Additionally, in the MLP classifier, we have the choice to insert our handcrafted features, giving the model more data to make its inference. These modifications to the ChangeFormer architecture can be seen in Figure 5.4.

5.2.1 Datasets and Preprocessing

To pre-train the encoder of ChangeFormer, we chose fluorescence microscopy images of cell nuclei. As a popular research area, cell nuclei imagery is widely available. Because of this, three different datasets were aggregated together. The first is a subset of the imagery provided by Kaggle’s 2018 Data Science Bowl[29] created by Schmidt et al.[13]. Their subset extracted only the fluorescence microscopy imagery and contains 21,636 instances of cell nuclei. The second dataset from Gunesli et al.[30] is composed of 3,329 liver cancer cell nuclei. The last dataset provided by Taschner-Mandl et al.[31] contains 7,813 cell nuclei from various tissue types. To prepare for training, each dataset was scaled to an appropriate resolution to match that of the AI-NP imagery. Next, all imagery was tiled into 256x256 images with 95% of empty tiles thrown out. Some were kept because the *after* images in the AI-NP imagery can be empty. Each dataset then had 15% of images withheld and combined for a test set. Combining the data in this manner kept the individual dataset distribution constant within the train and test set. The train and test sets are composed of 1,366 and 242 images respectively. Augmentations of rotations, flips, brightness jitter, and blur were incorporated during training, consistent with established augmentation schemes known to enhance the performance of deep neural models used on overhead imagery[32].

Input	Model	Precision	Recall	F1	IoU
Cells	Cell Encoder	0.952	0.956	0.954	0.912
AI-NP	Cell Encoder	0.878	0.596	0.710	0.550
	Cell Encoder, Refined on AI-NP	0.880	0.866	0.873	0.775
	Prev. StarDist AI-NP Segmentor	0.880	0.858	0.869	0.691

Table 5.1: Results of AI-NP segmentation through transfer learning on cell nuclei

5.2.2 Experimental Results

In order to utilize ChangeFormer for reaction classification, we needed to pre-train its encoder on cell nuclei, our analog for AI-NP imagery. These results can be seen in the first row of Table 5.1. Row two shows the results of running the AI-NP images through the unrefined cell encoder to serve as a baseline. F1 score of 0.710 and IoU of 0.55 once again allude that fluorescence microscopy images of cell nuclei are a nice substitute for optical AI-NP imagery. We proceeded to refine the encoder on our AI-NP imagery. The refining process demonstrated significant performance improvements in Recall, F1, and IoU scores. Additionally, when comparing it to our previous AI-NP segmentation model with a StarDist[13] backbone, we see performance increases again, especially in IoU as this encoder no longer predicts only star-convex polygons. We can now use this higher-performing AI-NP segmentor to regenerate more accurate handcrafted features. To do this, the new segmentor is inserted into the background subtraction methodology devised in chapter 2 and outlined in Figure 2.10. These new and more accurate features as well as the older features will be utilized in our classification models.

Our next experiment utilizes this trained encoder and the features it generates. We insert the encoder back into our modified ChangeFormer model and run the additional pre-training step forming change maps. Finally, we lock our encoder and train for reaction classification utilizing either the imagery or the imagery and handcrafted features on our two different pre-trained models. The results of this and a comparison

Model	Input	Pre-training Process	Accuracy	F1 Sinter	F1 Combination	F1 Spallation
ChangeFormer	Images	Segmentation	0.908	0.940	0.789	0.920
		Change Maps	0.863	0.895	0.695	0.886
	Images & Features	Segmentation	0.911	0.944	0.838	0.902
		Change Maps	0.875	0.893	0.769	0.896
	Images & Features V2	Segmentation	0.920	0.947	0.849	0.916
		Change Maps	0.878	0.897	0.737	0.906
Ensemble, Prev. Best	Features	NA	0.894	0.907	0.675	0.952

Table 5.2: ChangeFormer Based Reaction Classification Results

to the previous best ensemble-based three-class model can be seen in Table 5.2.2.

The pre-training step generating change maps is shown to be ineffective. A binary change map alone cannot accurately classify an image reaction pair because different classes can have identical binary change maps, an example of which can be seen in Figure 5.5. Even though performing change detection to form change maps is a similar task to reaction classification, the process loses too much information and steers a reaction classification model in the wrong direction.

ChangeFormer increased our overall accuracy in reaction classification with the best input being both the improved handcrafted features and imagery. The F1 score for spallations being highest with only feature input is a strong indication that the handcrafted features excel at identifying spallations. Another signal of this is the large performance increase in spallation F1 when moving from the old handcrafted features to improved features. Because of this, one would expect models with both inputs of handcrafted features and imagery to outperform or tie models with only features for input. The fact that we do not observe this signals possible data fusion issues where the model struggles to effectively leverage both inputs simultaneously. By adding the imagery to the input, we were able to greatly increase our F1 scores for sinters and combinations. These scores increase again when both handcrafted features and imagery are used and are less affected by the data fusion issues because the handcrafted features are not as effective on these two classes.




Figure 5.5: Two different classes yielding identical change maps

Our last set of experiments with ChangeFormer explores its adaptations and pre-training steps in a two-stage hierarchical architecture. Stage one separates the spallations from the sinters and combinations, and stage two separates the sinters and combinations. This hierarchical method was shown to be effective in chapter 4 because of the significant class overlap between the sintering and combination classes. These results are shown in Table 5.2.2.

Once again, we see performance degradation from the change map pre-training step, which further supports the claim that change maps alone are not enough information for classification and steer the model in the wrong direction. For step one, the ChangeFormer-based model was unable to outperform our previous best. This is most likely due to the hypothesized data fusion issues that partially stems from an insufficient amount of training data for such a large model despite our pre-training techniques. Unlike in the three-class model, step one shows decreased performance when utilizing features V2 over the original. Making matters more complicated, step two shows the best input is both the improved handcrafted features and images. This could signal that the more accurate features aid in describing the sinters and combinations more than the spallations. When moving from the previous ensemble

Model	Input	Pre-training Process	Accuracy	F1 Sinter	F1 Combination	F1 Spallation
ChangeFormer	Images	Segmentation	0.962		0.970	0.947
		Change Maps	0.935		0.947	0.915
	Images & Features	Segmentation	0.968		0.975	0.957
		Change Maps	0.932		0.945	0.909
	Images & Features V2	Segmentation	0.962		0.970	0.947
		Change Maps	0.955		0.965	0.938
Ensemble, Prev. Best	Features	NA	0.984		0.979	0.987
ChangeFormer	Images	Segmentation	0.969	0.980	0.923	[HTML]C0C0C0
		Change Maps	0.917	0.946	0.817	[HTML]C0C0C0
	Images & Features	Segmentation	0.974	0.984	0.930	[HTML]C0C0C0
		Change Maps	0.917	0.948	0.790	[HTML]C0C0C0
	Images & Features V2	Segmentation	0.979	0.987	0.941	[HTML]C0C0C0
		Change Maps	0.931	0.957	0.833	[HTML]C0C0C0
Ensemble, Prev. Best	Features	NA	0.828	0.890	0.612	[HTML]C0C0C0

Table 5.3: Hierarchical ChangeFormer Based Reaction Classification Results

models to ChangeFormer, we see significant gains in sinter and combination F1 performance. These performance gains are greater in the hierarchical architecture than in the three-class architecture.

Chapter 6

Autonomous SEM Imagery

Up until now, all imagery has been imaged through 100x magnification optical microscopy. We have shown that this is sufficient for effective reaction classification, however, there are details of the AI-NP imagery that cannot be captured this way. SEM imagery offers much higher magnification which allows for better texture analysis and can image smaller particle clusters or even single particles. An example of a single particle reaction that cannot be resolved with optical microscopy can be seen in Figure 6.1. Currently, the autonomous imaging acquisition system only works on optical imaging. This is because it used ImageJ's included particle analysis plugin (which could be upgraded to one of the AI-NP segmentors) to detect cluster centers for laser targeting. The SEM workflow has no equivalent and thus a method for autonomous targeting was desired. This targeter would take in an image like the one shown in Figure 6.2 and export the coordinates of all the cluster centers. These coordinates could then be translated into real-world locations so that the piezoelectric stage can align clusters for laser irradiation. To create this targeter, a dataset composed of 18 different images from various slides was collected.

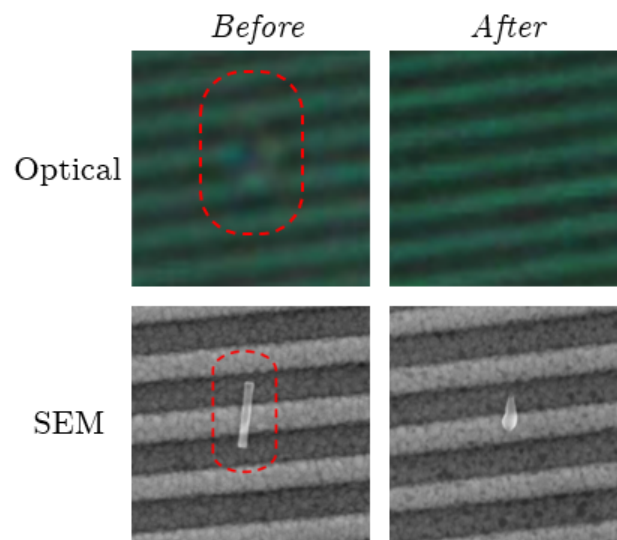


Figure 6.1: Single Particle Reaction Image Modality Comparison




Figure 6.2: SEM Example Image

6.1 Segment Anything

Segment Anything[33] is a transformer-based segmentation model created by Meta AI. Their publicly available pre-trained models were trained on 11 million high-resolution images which within there are 1.1 billion segmentation masks. The large volume of highly diverse training samples is why they claim the model is capable of segmenting anything. If this model can segment our SEM imagery we could then use the masks to extract cluster centers. Segment Anything performs its segmentation in two steps. It first runs the input image through an image encoder to create its embedding. This embedding is then run through a decoder with additional information, an encoded prompt, to create its masks. Segment Anything requires prompts through two different methods, points and or bounding boxes. A user can add points on an image denoting foreground and or background and Segment Anything will then segment the image with this added context. Additionally, users can supply bounding boxes signifying foreground. Bounding boxes and points may be used in conjunction. This use case is not very helpful for our autonomous targeted, as we do not have any information about where the clusters are to supply the model with as a prompt.

Segment Anything does have a prompting formula for autonomous use, a grid of foreground points. By laying a grid of foreground points equally over the image you sample the majority of the image and can assume you hit all targets. The issue however is how many points in the grid do not lie on a cluster target. This leads to erroneous mask regions due to random model interpretations or background slide imperfections. Figure 6.3 shows how applying a grid of points over an image yields many erroneous mask regions.




Figure 6.3: Segment Anything Prediction of Figure 6.2

6.2 Mask Filtering

For mask predictions from Segment Anything to be useful to us, the false positive regions need to be removed. The example in 6.3 shows mask regions that are long and easily filterable through eccentricity thresholding. However, the regions shown in Figure 6.4 cannot be filtered in this way. The false positive regions shown are much more circular and therefore another method for mask filtering must be determined. Because this process is being developed for an autonomous targeter we are much more concerned with the erroneous mask regions rather than any clusters Segment Anything has missed. The laser irradiation process and stage movement is time-consuming and therefore missing a cluster is much less important than avoiding irradiating the background and adding a ‘fake’ image to the dataset.




Figure 6.4: Segment Anything Prediction with Grating Imperfections

6.2.1 Texture

Now that we are able to make out the minute details in the clusters, it is easy to see the texture differences between the background and the clusters. For our texture descriptor, we chose the local binary pattern[34] (LBP) with rotational invariance. The LBP texture descriptor analyzes each pixel within a mask region and its surrounding pixels to see if they are of greater or less intensity than itself. Each pixel can then be assigned to a pattern depending on which of its neighbors is brighter or not. A histogram of the number of times each pattern was encountered can then be made to describe the texture of the area. We utilized a radius of three pixels at 8 equidistant angles. The masks were eroded to remove three pixels of depth all around the perimeter of each individual region. This ensures that entries into the LBP histogram can only look within the region area.

To begin, we calculated the LBP over an image with grating blemishes. This image is shown in Figure 6.5 where the pixel intensity represents which pattern (1-26) it corresponds to. In the LBP image, the cluster regions are still highly visible whereas grating blemishes blend in the background completely. This is a great signal that the LBP is highly distinctive when used to compare the two region types.

To create a classifier we collected histograms from seven cluster regions, true




Figure 6.5: LBP Image Example

positives (TPs), and seven regions that should be filtered out, false positives (FPs) to use as exemplars. The regions chosen were randomly selected with size brackets. For both the TPs and FPs four regions were selected from pixels sizes $60 \geq r < 1000$, two regions from $1000 \geq r < 4000$, and the last region from $5500 \geq r < 6500$. These size brackets were chosen to reflect the size distribution of our dataset, Figure 6.6. The histograms of all the exemplars can be seen in Figure 6.7, the last bin has been removed for visual clarity. The TP and FP histograms are highly characterized by their differences in the bins at edges and centers. The edge bins of an LBP histogram represent the flat unchanging patterns whereas the center bins represent edge and corner patterns. This is highly consistent with what we see in the SEM imagery as real cluster regions have many edges and corners due to the individual particles, and the FP regions are mostly uniform in intensity with the exception of the grating lines which have little area per region.

6.2.2 Experimental Results

To classify an individual region we compared it against the exemplars. To compare individual histograms we utilized Euclidean distance. However, there were a few options in how this was done. The first was how much of the two histograms to compare. We could simply compare the entire histogram or we could compare only




Figure 6.6: Size Histogram for All Mask Regions

the regions where they typically differ. Thus, we chose three different subsets of the histogram: the entire histogram, the edges (bins 1-5 and 21-25), or the center (bins 6-20). We also had a choice of which exemplars to compare the regions to. We could compare each region to all exemplars or compare them based on size. Small regions are compared to the small exemplars and small-medium regions (> 750 & < 1500) are compared to both small and medium exemplars. Medium regions to medium exemplars and lastly medium-large and large regions (> 3500) to both the medium exemplars and the single large exemplar. After the exemplars to compare against are chosen, the Euclidean distance between the target region and each FP exemplar is summed and recorded. Similarly, the Euclidean distances between the target region and each TP exemplar are summed separately. The target region is labeled as whichever sum is lesser.

Table 6.1 shows the results from the experimentation over the different exemplar selection techniques and bin subsets. Again, we are highly concerned with minimizing the false positives to ensure we do not add false clusters to our dataset. Because of

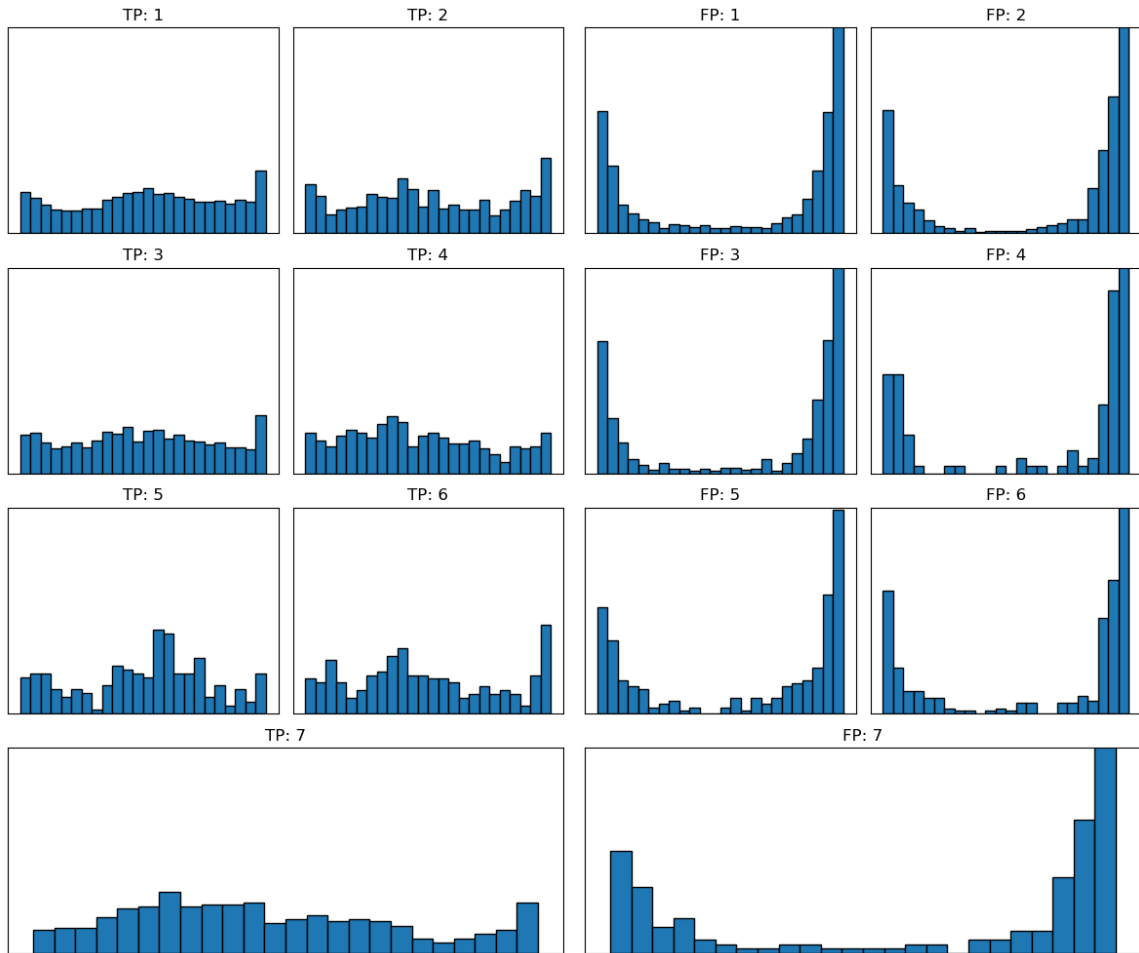


Figure 6.7: Histogram of All Exemplars: TPs(left) FPs(right)

Exemplar Selection	Bins	Accuracy	Precision	Recall	F1
All	All	0.982	0.982	0.997	0.990
	Center	0.968	0.992	0.971	0.982
	Edges	0.984	0.982	1.0	0.991
Size	All	0.984	0.985	0.997	0.991
	Center	0.961	1.0	0.956	0.977
	Edges	0.984	0.982	1.0	0.991

Table 6.1: Texture Based Filtering Results

this, we happily accept lower performance metrics in overall accuracy, recall, and F1 to use the model with perfect precision. Interestingly, utilizing only the center bins when taking the Euclidean distance gives the best precision in both exemplar selection methods. While the best model does use size to select exemplars the small difference in performance and our small dataset means we cannot say for certain if this is the overall superior method. With this texture-based algorithm for filtering out false Segment Anything mask regions, we can confidentially image and target clusters for laser irradiation autonomously.

Chapter 7

Summary and concluding remarks

We have presented a processing pipeline for analyzing nanoenergetic material reaction pairs. The pipeline performs background subtraction, leaving behind only the reacted area within the *before* and *after* reaction pair. This preprocessing step is valuable for accurately characterizing the reaction that occurred. Furthermore, the presented pipeline utilizes techniques that are effective even on smaller datasets, making it ideal for nanoenergetics. We introduce a set of feature extractors that are applicable to the background subtracted imagery of Al-NPs along with the corresponding methodology to compare the pre- and post-reaction measurements to enable reaction classification. Using the presented features, we showed significant overlap between the classes of Sintering and Combination reactions. To address this, we proposed a hierarchical classification approach that proved to be more effective than the three-class models. We not only demonstrated the effectiveness of the features in reaction classification but also identified which specific features are particularly useful for distinguishing between the different reaction classes.

Moving to deep machine learning we have presented an alteration to a popular

state-of-the-art change detection model to enable reaction classification. The adaptation is modular and allows for supplemental feature additions. We have also illustrated fluorescence microscopy imagery of cell nuclei to be an effective analog for difficult-to-obtain Al-NP imagery. Using this image analog in pre-training allowed effective training of a large transformer-based model even with minimal target data. This model may also be repurposed to improve the autonomous processing pipeline for analyzing nanoenergetic material reaction pairs and for more accurate feature extraction. Conversely to pre-training with cell nuclei, the generation of change maps as a pre-training step for reaction classification was shown to be ineffective and degrades performance. The handcrafted features included in the model input were shown to increase overall performance but reduce spallation performance due to suspected data fusion issues. We then reaffirmed that a hierarchical model configuration greatly aids in remedying the class overlap issue between the sinters and combinations.

These reaction classification models help us further our goal of reaction prediction. We aim to aid researchers in characterizing the heating rate needed to elicit the MDM given the composition, size, and shape of both an Al-NP or a cluster of Al-NPs. The autonomous reaction classifier we have built can serve as a labeler as we amass a large dataset based on a grid of variables including the aforementioned material properties, but also laser properties like power density, pulse duration, and wavelength. We may then use this grid of reaction labels and laser and cluster properties to search for a quantitative mapping between the properties and reaction labels.

Lastly, we began experimentation with SEM imagery. We show texture is a highly distinctive feature for separating Al-NP cluster masks from background grating. Through this, we illustrate a methodology for utilizing Segment Anything without human prompting through the filtering of mask regions. The process in its entirety allows for the autonomous imaging and laser irradiation of Al-NP material with SEM

imagery. This imagery can be collected in conjunction with optical imagery in the grid dataset. The higher magnification SEM imagery allows for the measurement of more physical parameters about Al-NPs to use in the search for a quantitative mapping between the reaction parameters and reaction label.

Bibliography

- [1] M. A. Trunov, M. Schoenitz, and E. L. Dreizin. Effect of polymorphic phase transformations in alumina layer on ignition of aluminium particles. *Combustion Theory and Modelling*, 10(4):603–623, 2006. eprint: <https://doi.org/10.1080/13647830600578506>.
- [2] A. Rai, K. Park, L. Zhou, and M. R. Zachariah. Understanding the mechanism of aluminium nanoparticle oxidation. *Combustion Theory and Modelling*, 10(5):843–859, 2006. eprint: <https://doi.org/10.1080/13647830600800686>.
- [3] V. I. Levitas, B. W. Asay, S. F. Son, and M. Pantoya. Melt dispersion mechanism for fast reaction of nanothermites. *Applied Physics Letters*, 89(7):071909, 2006.
- [4] V. I. Levitas, M. L. Pantoya, and K. W. Watson. Melt-dispersion mechanism for fast reaction of aluminum particles: Extension for micron scale particles and fluorination. *Applied Physics Letters*, 92(20):201917, May 2008. ISSN: 0003-6951. eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.2936855/14395299/201917_1_online.pdf.
- [5] V. I. Levitas. Mechanochemical mechanism for reaction of aluminium nano- and micrometre-scale particles. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(2003):20120215, 2013.

eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2012.0215>.

- [6] D. S. Sundaram, P. Puri, and V. Yang. A general theory of ignition and combustion of nano- and micron-sized aluminum particles. *Combustion and Flame*, 169:94–109, 2016. ISSN: 0010-2180.
- [7] V. I. Levitas, B. W. Asay, S. F. Son, and M. Pantoya. Mechanochemical mechanism for fast reaction of metastable intermolecular composites based on dispersion of liquid metal. *Journal of Applied Physics*, 101(8):083524, Apr. 2007. ISSN: 0021-8979. eprint: https://pubs.aip.org/aip/jap/article-pdf/doi/10.1063/1.2720182/14981427/083524\1_online.pdf.
- [8] N. Zakiyyan, C. Mathai, J. McFarland, S. Gangopadhyay, and M. R. Maschmann. Spallation of Isolated Aluminum Nanoparticles by Rapid Photothermal Heating. *ACS Applied Materials & Interfaces*, 14(49):55277–55283, 2022. PMID: 36445833. eprint: <https://doi.org/10.1021/acsami.2c18678>. URL: <https://doi.org/10.1021/acsami.2c18678>.
- [9] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri. NIH Image to ImageJ: 25 years of image analysis. *Nature Methods*, 9(7):671–675, July 2012. ISSN: 1548-7105. URL: <https://doi.org/10.1038/nmeth.2089>.
- [10] J. Kapur, P. Sahoo, and A. Wong. A new method for gray-level picture thresholding using the entropy of the histogram. *Computer Vision, Graphics, and Image Processing*, 29(3):273–285, 1985. ISSN: 0734-189X. URL: <https://www.sciencedirect.com/science/article/pii/0734189X85901252>.
- [11] P. Thevenaz, U. Ruttimann, and M. Unser. A pyramid approach to subpixel registration based on intensity. *IEEE Transactions on Image Processing*, 7(1):27–41, 1998.

- [12] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, 1987. ISSN: 0734-189X. URL: <https://www.sciencedirect.com/science/article/pii/S0734189X8780186X>.
- [13] U. Schmidt, M. Weigert, C. Broaddus, and G. Myers. Cell Detection with Star-convex Polygons. *CoRR*, abs/1806.03535, 2018. arXiv: [1806.03535](https://arxiv.org/abs/1806.03535). URL: <http://arxiv.org/abs/1806.03535>.
- [14] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *CoRR*, abs/1505.04597, 2015. arXiv: [1505.04597](https://arxiv.org/abs/1505.04597). URL: <http://arxiv.org/abs/1505.04597>.
- [15] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017. arXiv: [1703.06870](https://arxiv.org/abs/1703.06870). URL: <http://arxiv.org/abs/1703.06870>.
- [16] M.-K. Hu. Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, 8(2):179–187, 1962.
- [17] W. LI, D. ZHANG, and Z. XU. PALMPRINT IDENTIFICATION BY FOURIER TRANSFORM. *International Journal of Pattern Recognition and Artificial Intelligence*, 16(04):417–432, 2002. eprint: <https://doi.org/10.1142/S0218001402001757>. URL: <https://doi.org/10.1142/S0218001402001757>.
- [18] R. M. Haralick, K. Shanmugam, and I. Dinstein. Textural Features for Image Classification. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3(6):610–621, 1973.

- [19] L. van der Maaten and G. Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. URL: <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- [20] L. Breiman. Random Forests. *Machine Learning*, 45(1):5–32, Oct. 2001. ISSN: 1573-0565. URL: <https://doi.org/10.1023/A:1010933404324>.
- [21] Y. Freund and R. E. Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. ISSN: 0022-0000. URL: <https://www.sciencedirect.com/science/article/pii/S002200009791504X>.
- [22] T. Hastie, S. Rosset, J. Zhu, and H. Zou. Multi-class adaboost. *Statistics and its Interface*, 2(3):349–360, 2009.
- [23] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001. URL: <https://doi.org/10.1214/aos/1013203451>.
- [24] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, Aug. 1996. ISSN: 1573-0565. URL: <https://doi.org/10.1007/BF00058655>.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [26] W. G. C. Bandara and V. M. Patel. *A Transformer-Based Siamese Network for Change Detection*. 2022. arXiv: [2201.01293](https://arxiv.org/abs/2201.01293) [cs.CV].

- [27] H. Chen and Z. Shi. A Spatial-Temporal Attention-Based Method and a New Dataset for Remote Sensing Image Change Detection. *Remote Sensing*, 12(10):1662, May 2020. ISSN: 2072-4292. URL: <http://dx.doi.org/10.3390/rs12101662>.
- [28] C. Zhang, P. Yue, D. Tapete, L. Jiang, B. Shangguan, L. Huang, and G. Liu. A deeply supervised image fusion network for change detection in high resolution bi-temporal remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:183–200, 2020. ISSN: 0924-2716. URL: <https://www.sciencedirect.com/science/article/pii/S0924271620301532>.
- [29] J. C. Caicedo, A. Goodman, K. W. Karhohs, B. A. Cimini, J. Ackerman, M. Haghighi, C. Heng, T. Becker, M. Doan, C. McQuin, M. Rohban, S. Singh, and A. E. Carpenter. Nucleus segmentation across imaging experiments: the 2018 Data Science Bowl. *Nature Methods*, 16(12):1247–1253, Dec. 2019. ISSN: 1548-7105. URL: <https://doi.org/10.1038/s41592-019-0612-7>.
- [30] G. N. Gunesli, C. Sokmensuer, and C. Gunduz-Demir. AttentionBoost: Learning What to Attend for Gland Segmentation in Histopathological Images by Boosting Fully Convolutional Networks. *IEEE Transactions on Medical Imaging*, 39(12):4262–4273, 2020.
- [31] S. Taschner-Mandl, I. M. Ambros, P. F. Ambros, K. Beiske, A. Hanbury, W. Doerr, T. Weiss, M. Berneder, M. Ambros, E. Bozsaky, and F. Kromp. *An annotated fluorescence image dataset for training nuclear segmentation methods*. <https://www.ebi.ac.uk/biostudies/bioimages/studies/S-BSST265>. 2020.
- [32] G. J. Scott, M. R. England, W. A. Starms, R. A. Marcum, and C. H. Davis. Training Deep Convolutional Neural Networks for Land-Cover Classification

of High-Resolution Imagery. *IEEE Geoscience and Remote Sensing Letters*, 14(4):549–553, 2017.

- [33] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick. Segment Anything. *arXiv:2304.02643*, 2023.
- [34] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, 2002.

VITA

Tim Gaines was born and raised in St. Louis, Missouri. He graduated from the Missouri University of Science and Technology in 2021 with a Bachelors of Science in both Electrical Engineering and Computer Science. During his undergraduate career he completed 3 internships at bioMérieux and 1 at Clairvoyant Networks. Additionally during the school year he worked at the Applied Microwave Nondestructive Testing Laboratory. Right after completing his undergraduate degrees he enrolled into graduate school and began his research into computer vision with Dr. Scott at the Center for Geospatial Intelligence. He then earned his Masters of Science in Computer Science from the University of Missouri - Columbia in 2023. Following the completion of his MS he moved to Portland, Oregon to be with his wife while she completes her residency. Tim is currently in the interviewing process for embedded software engineer positions.