

TEXTURE-DRIVEN IMAGE CLUSTERING IN LASER POWDER BED FUSION

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science

by

ALEXANDER H. GROEGER
B.S.C.E., Wright State University, 2020
B.S., Wright State University, 2020

2021
Wright State University

WRIGHT STATE UNIVERSITY

GRADUATE SCHOOL

November 22, 2021

I HEREBY RECOMMEND THAT THE THESIS PREPARED UNDER MY SUPERVISION BY Alexander H. Groeger ENTITLED Texture-Driven Image Clustering in Laser Powder Bed Fusion BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF Master of Science.

Tanvi Banerjee, Ph.D.
Thesis Director

Michael L. Raymer, Ph.D.
Chair, Department of Computer
Science and Engineering

Committee on Final Examination

Tanvi Banerjee, Ph.D.

John Middendorf, Ph.D.

Thomas Wischgoll, Ph.D.

Barry Milligan, Ph.D.
Vice Provost for Academic Affairs
Dean of the Graduate School

ABSTRACT

Groeger, Alexander H. M.S., Department of Computer Science and Engineering, Wright State University, 2021. *Texture-Driven Image Clustering in Laser Powder Bed Fusion*.

The additive manufacturing (AM) field is striving to identify anomalies in laser powder bed fusion (LPBF) using multi-sensor in-process monitoring paired with machine learning (ML). In-process monitoring can reveal the presence of anomalies but creating a ML classifier requires labeled data. The present work approaches this problem by printing hundreds of Inconel-718 coupons with different processing parameters to capture a wide range of process monitoring imagery with multiple sensor types. Afterwards, the process monitoring images are encoded into feature vectors and clustered to isolate groups in each sensor modality. Four texture representations were learned by training two convolutional neural network texture classifiers on two general texture datasets for clustering comparison. The results demonstrate unsupervised texture-driven clustering can isolate roughness categories and process anomalies in each sensor modality. These groups can be labeled by a field expert and potentially be used for defect characterization in process monitoring.

Contents

1	Introduction	1
1.1	Motivation and Objective	1
1.2	Organization	4
2	Background	5
2.1	Laser Powder Bed Fusion	5
2.2	Image Clustering	10
2.3	Texture Analysis Using Image Preprocessing	11
3	Experiments	13
3.1	Datasets	13
3.1.1	Texture Datasets	14
3.1.2	LPBF In-situ Sensor Data	19
3.2	Models	26
3.2.1	TextureCNN	27
3.2.2	WaveletCNN	28
3.3	Fast Density Clustering	29
3.4	UMAP Dimensionality Reduction	30
4	Results and Discussion	32
4.1	Texture Datasets	34
4.1.1	Kylberg	35
4.1.2	ALOT	53
4.2	LPBF In-situ Analysis	68
4.2.1	Recoat Post Spread	69
4.2.2	Post Melt	102
4.2.3	Thermal Tomography	139
4.2.4	LWIR	175
4.2.5	Spatter	203
5	Conclusion	234
5.1	Data Collection and Preprocessing	234

5.1.1	Texture Datasets	234
5.1.2	LPBF Datasets	234
5.2	Texture-Driven Strategy	236
5.3	Contributions	237
5.3.1	Kylberg	237
5.3.2	ALOT	238
5.3.3	LPBF Datasets	238
5.3.4	Recoat Post-Spread	238
5.3.5	Recoat Post-Melt	240
5.3.6	Thermal Tomography	242
5.3.7	LWIR	243
5.3.8	Spatter	244
5.3.9	Model Evaluation	245
5.4	Final Thoughts	246
6	Future Work	248
6.1	Modifications	248
6.1.1	Data Collection and Preprocessing	248
6.1.2	Model Architecture	251
6.1.3	Clustering	252
6.2	LPBF Tailored Approach	252
	Bibliography	255

List of Figures

1.1	Illustration of machining vs additive manufacturing.	1
1.2	Flow diagram of an experiment. All 4 experiments use a unique combination of texture training set and deep learning model that produces 5 texture representations for analysis.	3
2.1	Diagram of LPBF machine accurate to OpenAdditive’s configuration.	6
2.2	Illustration of continuous hatching vs striped hatching.	9
2.3	Illustration of where laser surface parameters are used.	10
3.1	Kylberg samples from each class.	15
3.2	Preprocessing steps for Kylberg.	16
3.3	ALOT samples from each class.	17
3.4	Preprocessing steps for ALOT.	18
3.5	Arrangement of cameras above and coupons within the build plane.	20
3.6	Examples of in-situ recoat images. The angled horizontal black line in both images is residue after preprocessing and laser damage to the sensor.	21
3.7	Example of thermal tomography for a single layer. The center left coupon is experiencing washout.	23
3.8	Example of peak temperatures from LWIR for a single layer. The hatch angle can be observed by the lines in the coupons.	24
3.9	Example of compiled spatter for a single layer. Notice the coupon tails. These are hot emissions and soot being pulled away by the airflow.	26
3.10	Our TextureCNN architecture.	28
3.11	Our WaveletCNN architecture.	29
3.12	UMAP Visualization of MNIST flattened image vectors parameterized with a minimum distance of 0.1 and 15 nearest neighbors. The points are colored by their class label. The homogeneity of color in each cluster indicates most samples in tight proximity share the same semantic label which we view as a desirable clustering.	31
4.1	Plot of the average activation in each feature dimension from Kylberg for both models trained on Kylberg. Lines indicate the range of the activations.	35

4.2	Plot of the average number of clusters Kylberg breaks into at different noise thresholds for two experiments. The range over 5 runs is indicated by a line on each bar.	36
4.3	Kylberg cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.2.	37
4.4	Normalized images sampled from each cluster in the Kylberg Kylberg-TCNN feature space.	41
4.5	UMAP visualization of the Kylberg Kylberg-TCNN feature space colored by class. Minimum distance of 0.25 and 300 neighbors were used.	43
4.6	UMAP visualization of the Kylberg Kylberg-TCNN feature space colored by cluster number. Minimum distance of 0.25 and 300 neighbors were used.	44
4.7	Kylberg cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.2.	45
4.8	Normalized images sampled from each cluster in the Kylberg Kylberg-WCNN feature space.	49
4.9	UMAP visualization of the Kylberg Kylberg-WCNN feature space. Minimum distance of 0.25 and 300 neighbors were used.	50
4.10	UMAP visualization of the Kylberg Kylberg-WCNN feature space. Minimum distance of 0.25 and 300 neighbors were used.	51
4.11	Plot of the average activation in each feature dimension from ALOT for both models trained on ALOT. Lines indicate the range of the activations.	53
4.12	Plot of the average number of clusters ALOT breaks into at different noise thresholds for two experiments. The range over 5 runs is indicated by a line on each bar.	54
4.13	ALOT cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.01.	55
4.14	Normalized images sampled from each cluster in the ALOT ALOT-TCNN feature space.	59
4.15	UMAP visualization of the ALOT ALOT-TCNN feature space colored by cluster. Minimum distance of 0.10 and 15 neighbors were used.	60
4.16	ALOT cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.01.	61
4.17	Normalized images sampled from each cluster in the ALOT ALOT-WCNN feature space.	65
4.18	UMAP visualization of the ALOT ALOT-WCNN feature space. Minimum distance of 0.10 and 15 neighbors were used.	66
4.19	Plot of the average number of clusters the post-spread data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.	69
4.20	Recoat post-spread cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.	70
4.21	Normalized images sampled from each cluster in the recoat post-spread Kylberg-TCNN feature space. We describe the textural roughness of the powder spread followed by textural descriptions of other defects.	75

4.22	UMAP visualization of the Recoat Post-Spread Kylberg-TCNN feature space colored by cluster number. Minimum distance of 0 and 50 neighbors were used.	77
4.23	Normalized images sampled from severe defects cluster in the recoat post-spread Kylberg-TCNN feature space UMAP projection.	78
4.24	UMAP visualization of the Recoat Post-Spread Kylberg-TCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.	79
4.25	Recoat post-spread cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.	81
4.26	Normalized images sampled from each cluster in the recoat post-spread Kylberg-WCNN feature space.	84
4.27	UMAP visualization of the Recoat Post-Spread Kylberg-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	85
4.28	UMAP visualization of the Recoat Post-Spread Kylberg-WCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.	86
4.29	Recoat post-spread cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.	88
4.30	Normalized images sampled from each cluster in the recoat post-spread ALOT-TCNN feature space.	91
4.31	UMAP visualization of the Recoat Post-Spread ALOT-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	92
4.32	UMAP visualization of the Recoat Post-Spread ALOT-TCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.	93
4.33	Recoat post-spread cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.	94
4.34	Normalized images sampled from each cluster in the recoat post-spread ALOT-WCNN feature space.	97
4.35	UMAP visualization of the Recoat Post-Spread ALOT-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	98
4.36	UMAP visualization of the Recoat Post-Spread ALOT-WCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.	99
4.37	Plot of the average number of clusters the post-melt data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.	102
4.38	Recoat post-melt cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.	103
4.39	Normalized images sampled from each cluster in the recoat post-melt Kylberg-TCNN feature space.	106
4.40	UMAP visualization of the Recoat Post-Melt Kylberg-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	108

4.41	UMAP visualization of the Recoat Post-Melt Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	110
4.42	Recoat post-melt cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.	112
4.43	Normalized images sampled from each cluster in the recoat post-melt Kylberg-WCNN feature space.	116
4.44	UMAP visualization of the Recoat Post-Melt Kylberg-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	118
4.45	UMAP visualization of the Recoat Post-Melt Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	119
4.46	Recoat post-melt cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.	120
4.47	Normalized images sampled from each cluster in the recoat post-melt ALOT-TCNN feature space.	125
4.48	UMAP visualization of the Recoat Post-Melt ALOT-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	127
4.49	UMAP visualization of the Recoat Post-Melt ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	129
4.50	Recoat post-melt cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.	130
4.51	Normalized images sampled from each cluster in the recoat post-melt ALOT-WCNN feature space.	133
4.52	UMAP visualization of the Recoat Post-Melt ALOT-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.	134
4.53	UMAP visualization of the Recoat Post-Melt ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	136
4.54	Plot of the average number of clusters the thermal tomography data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.	140
4.55	Thermal tomography cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.	141
4.56	Normalized images sampled from each cluster in the thermal tomography Kylberg-TCNN feature space.	146
4.57	UMAP visualization of the Thermal Tomography Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	148
4.58	UMAP visualization of the Thermal Tomography Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	150
4.59	Thermal tomography cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.	151

4.60	Normalized images sampled from each cluster in the thermal tomography Kylberg-WCNN feature space.	155
4.61	UMAP visualization of the Thermal Tomography Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	156
4.62	Normalized images sampled from the spatter regions in the thermal tomography Kylberg-WCNN feature space UMAP projection.	157
4.63	UMAP visualization of the Thermal Tomography Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	158
4.64	Thermal tomography cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.	159
4.65	Normalized images sampled from each cluster in the thermal tomography ALOT-TCNN feature space.	163
4.66	UMAP visualization of the Thermal Tomography ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	164
4.67	UMAP visualization of the Thermal Tomography ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	167
4.68	Thermal tomography cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.	168
4.69	Normalized images sampled from each cluster in the thermal tomography ALOT-WCNN feature space.	170
4.70	UMAP visualization of the Thermal Tomography ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	171
4.71	UMAP visualization of the Thermal Tomography ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	172
4.72	Plot of the average number of clusters the LWIR data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.	175
4.73	LWIR cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.	176
4.74	Normalized images sampled from each cluster in the LWIR Kylberg-TCNN feature space.	179
4.75	UMAP visualization of the LWIR Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	180
4.76	UMAP visualization of the LWIR Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	182
4.77	LWIR cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.	183
4.78	Normalized images sampled from each cluster in the LWIR Kylberg-WCNN feature space.	186

4.79	UMAP visualization of the LWIR Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	188
4.80	UMAP visualization of the LWIR Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	189
4.81	LWIR cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.	190
4.82	Normalized images sampled from each cluster in the LWIR ALOT-TCNN feature space.	192
4.83	UMAP visualization of the LWIR ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	194
4.84	Normalized images sampled from the spatter region in the LWIR ALOT-TCNN feature space UMAP projection.	195
4.85	UMAP visualization of the LWIR ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	196
4.86	LWIR cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.	197
4.87	Normalized images sampled from each cluster in the LWIR ALOT-WCNN feature space.	198
4.88	UMAP visualization of the LWIR ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	199
4.88	UMAP visualization of the LWIR ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	200
4.89	Plot of the average number of clusters the spatter data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.	204
4.90	Spatter cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.	205
4.91	Normalized images sampled from each cluster in the spatter Kylberg-TCNN feature space.	209
4.92	UMAP visualization of the Spatter Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	211
4.93	UMAP visualization of the spatter Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	213
4.94	Spatter sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.	214
4.95	Normalized images sampled from each cluster in the spatter Kylberg-WCNN feature space.	216
4.96	UMAP visualization of the Spatter Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	217
4.97	UMAP visualization of the spatter Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.	219
4.98	Spatter cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.	220
4.99	Normalized images sampled from each cluster in the spatter ALOT-TCNN feature space.	222

4.100	UMAP visualization of the Spatter ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	223
4.101	Normalized images sampled from the spatter region in the Spatter ALOT-TCNN feature space UMAP projection.	224
4.102	UMAP visualization of the spatter ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used. . .	225
4.103	Spatter cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.	226
4.104	Normalized images sampled from each cluster in the spatter ALOT-WCNN feature space.	228
4.105	UMAP visualization of the Spatter ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.	229
4.106	UMAP visualization of the spatter ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used. . .	231

List of Tables

3.1	Pool of processing parameters. All 243 combinations of these parameters were built by switching parameter sets every so many layers across 9 coupons.	19
4.1	Classification accuracies of each experiment.	32
4.2	Table of LPBF phenomena seen in Kylberg-TCNN recoat post-spread clusters.	76
4.3	Table of surface categories seen in Kylberg-WCNN recoat post-melt clusters.	117
4.4	Table of process phenomenon seen in Kylberg-TCNN thermal tomography clusters.	147
4.5	Table of process phenomenon seen in Kylberg-WCNN LWIR clusters.	187
4.6	Table of process phenomenon seen in Kylberg-TCNN Spatter clusters.	209

List of Acronyms

Material Science

AM	Additive Manufacturing
LPBF	Laser Powder Bed Fusion
LWIR	Long Wave Infrared

Computer Science

ALOT	Amsterdam Library of Textures
CNN	Convolutional Neural Network
TCNN	Texture Convolutional Neural Network
WCNN	Wavelet Convolutional Neural Network
UMAP	Uniform Manifold Approximation and Projection
t-SNE	t-Distributed Stochastic Embedding

Other

2D	Two-Dimensional
3D	Three-Dimensional

Acknowledgments

I thank Jason Knapp with much enthusiasm for donating his time to advise me and direct the development of various texture-based approaches.

I am grateful to John Middendorf, Joe Walker, Matthew Larson, Chris Barrett, and many others at OpenAdditive for welcoming me into their workspace and including me in their operations for the benefit of my learning.

My advisers, Tanvi Banerjee and Joy Gockel, were a blessing for their patience, advising, and investment in my academic career. I am very thankful to have known them.

Finally, I thank my committee Tanvi Banerjee, John Middendorf, and Thomas Wischgoll for evaluating my thesis and defense presentation.


This project was funded through Award DE-EE0009097 at Wright State University by the Department of Energy, a summer grant from the Ohio Space Grant Consortium, and the Material Engineering department at Wright State University.

Dedicated to
my family and friends

Introduction

1.1 Motivation and Objective

The manufacturing industry has traditionally practiced subtractive manufacturing. One of the more common processes is machining. Parts are built by carving or cutting away material from a block or sheet to achieve a desired geometry. While effective, this approach produces a lot of wasted material, which elevates production costs. Additive manufacturing (AM) provides a solution to this by taking the opposite approach of building onto the part until a desired geometry is realized. This allows for more complex geometries while reducing material waste. AM typically achieves this advantage by slicing a part's geometry into very thin 2D cross sections, called layers, and building these layers on top of each other one at a time. Figure 1.1 depicts the difference between machining and additive manufacturing. There are many unique processes to AM, but the one we are particularly interested in is laser powder bed fusion (LPBF).



(a) Machining uses carving tools on bulk material to shape the final part.

(b) Additive Manufacturing builds successive layers to shape a part.

Figure 1.1: Illustration of machining vs additive manufacturing.

The LPBF process is continuously being optimized for consistency and reliability in part quality. There are many physical phenomena that arise from the interaction between the laser and powder that influence the quality of the fabricated part by the integration of pores. It is important to manufacturers that these phenomena are monitored and controlled within acceptable levels to guarantee the consistent production of high-quality parts with specific properties. This problem is addressed using process monitoring.

Most approaches to process monitoring use in-situ sensors to capture different aspects of the process. These sensors come in different modalities that capture distinct regions of the light spectrum using specific exposure times and frame rates to extract unique information from the process. Some examples of these modalities include visible recoat (post-spread and post-fusion imaging) and infrared video. The use of analytics in process monitoring can aid manufacturers in the development of processing parameters when working with new materials, defect mitigation strategies, and anomaly detection notification systems of which all save time and resources in production. One challenge with developing statistical-model analytics for LPBF data is generalization as they typically are tuned for a specific material and lighting with fixed camera configurations (i.e. gain and exposure). For more advanced analytics, machine learning can be utilized, but a majority of tasks require data labeled by a field expert, which in this work there are no labels nor is it known every phenomenon that can be captured by each sensor modality. This thesis work overcomes this limitation by using an unsupervised approach to reveal inter-data relationships.

The choice of method is image clustering where the task is to isolate groups of similar images that cluster together in a high-dimensional space. Since most in-situ monitoring data is recorded as images, this approach can be applied to all sensor modalities to isolate image patches containing similar features contributing toward similar phenomena. This approach requires little field expert knowledge and gives unbiased data-driven groupings.

The goal is to provide LPBF manufacturers with better understanding of their process monitoring data and how it can be segmented for anomalous phenomena. This problem is ap-

proached using texture analysis and deep learning without the need for labels. Texture is emphasized because bulk regions within a part’s layer geometry do not contain significant edges or shape information.

This thesis proposes an approach that addresses the problem of unsupervised analysis of LPBF data from several process monitoring sensor modalities by embedding images into a compact textural representation using deep learning. Figure 1.2 shows the process flow for our experiments.

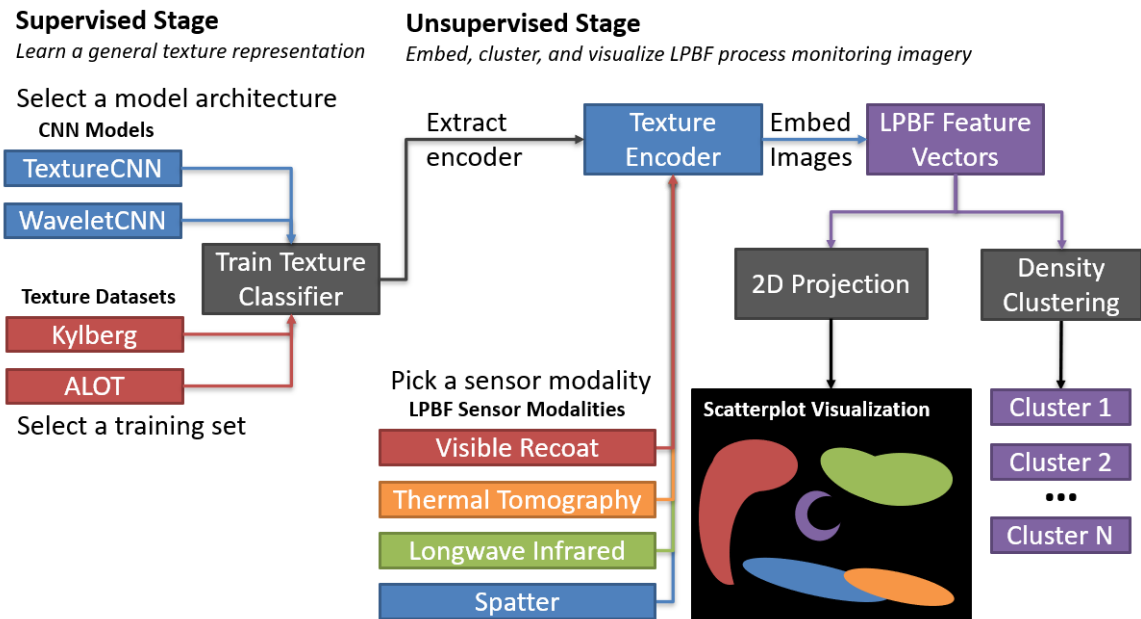


Figure 1.2: Flow diagram of an experiment. All 4 experiments use a unique combination of texture training set and deep learning model that produces 5 texture representations for analysis.

The contributions of this work are the following.

1. Four experiments are performed each of which train one of two deep learning models for texture classification on one of two texture datasets. The trained encoders are repurposed afterwards to extract and embed textural features from each modality of the LPBF data to obtain textural feature spaces.

2. Each sensor modality feature space is clustered using a hard density-based clustering algorithm and samples from each cluster are visualized to understand what LPBF process phenomena are being segmented in each sensor modality and why.
3. The arrangement of texture features and biases in texture distributions due to external factors are revealed through colored projections of those feature spaces from each sensor modality.

1.2 Organization

This thesis is organized as follows. In Chapter 2, we discuss the background of the AM LPBF technology we are studying in Section 2.1, an overview of image clustering in Section 2.2, and texture analysis in Section 2.3. In Chapter 3, we introduce the texture training data and LPBF data to be analyzed in Section 3.1, how we adopted and modified texture-based deep learning architectures in Section 3.2, the clustering method in Section 3.3, and how we visualize feature spaces in Section 3.4. In Chapter 4, we show our results and analyses of the texture data in Sections 4.1.1 (Kylberg) and 4.1.2 (ALOT), and the LPBF data in Sections 4.2.1 (Recoat post-spread), 4.2.2 (Recoat post-melt), 4.2.3 (Thermal tomography), 4.2.4 (LWIR), and 4.2.5 (Spatter). In Chapter 5, we summarize our conclusions. Finally, in Chapter 6, we discuss future work that can build off of this work.

Background

2.1 Laser Powder Bed Fusion

Laser Powder Bed Fusion (LPBF) is an additive manufacturing process where parts are constructed layer-by-layer by melting and fusing metallic powder with a laser. Several machine configurations exist, so we provide background information that pertains to the machine we used from OpenAdditive, LLC.

The process is held within a chamber to isolate its environment from the outside environment because welding in LPBF requires a low-oxygen environment. This is achieved by pumping the chamber with a non-reactive gas, typically argon, to drive out oxygen. Within the chamber, powder is managed by a build plate, feeder, recoater, and catch. See the machine setup illustration in Figure 2.1. The build plate is initially positioned at the same level of the build plane. When a layer is built, the build plate retracts leaving a gap for a thin layer of powder to be received. This displacement fixes the height of the build plane to avoid the need for refocusing the laser after each layer and prevents the part from obstructing the recoater while it spreads a new layer of powder. The powder feeder pushes fresh powder on a piston-actuated plate up above the build plane. The powder recoater swipes over the powder feed plate using a soft silicon blade to collect powder above the build plane. As the recoater moves, it deposits powder uniformly along the build plate to fill the gap between the build plate and plane. Excess powder is pushed into the overflow powder catch for recycling. Then a class 4 laser welds the freshly spread powder onto the

build plate or prior layer of the part. This process is repeated until the part is complete. The welding process produces soot which can create problems. Flying hot particles landing and fusing onto a part can cause build issues or even build failure. The laser performance can be degraded either from the plume and smoke obstructing the laser beam or smoke rising and covering the laser window. These problems are mitigated using an airflow across the build layer to catch and drag the soot away. The soot is collected by an air filter which circulates clean argon back into the airflow. As a part is built up, the build plate retracts using a piston to offset the new layer. Process monitoring cameras capture data on the build layer throughout the process until the part is completed.




Figure 2.1: Diagram of LPBF machine accurate to OpenAdditive’s configuration.

One of the challenges in LPBF is understanding how to control part quality. Processing monitoring is often necessary to accomplish this. Most LPBF machines come equipped with a visible recoat camera to capture post-spread and post-melt images of the build plate while some systems additionally supply a thermal camera to capture weld emissions. These data can be analyzed to find anomalous activities, discover how they affect

the quality of the part, and learn what elements of the process cause these phenomena and how to control them. A recent research review on in-situ monitoring [11] reports much research has been done on analyzing process monitoring data for part defects using machine learning algorithms.

AddiGuru [2] developed an image segmentation approach to locate defects in LPBF post-spread images. These defects include recoat streaking, recoat hopping, and material protrusion whose characterization from powder spreads has been studied in [25]. Recoat streaking and hopping can be detected by finding parallel or perpendicular lines to the recoater movement direction in the powder. Material protrusion can be detected by finding melted material poking up through the powder layer. AddiGuru is a business selling products, therefore the details of their algorithms remain unknown to the public.

Scime et al. [26] developed a layer-wise pixel-wise post-spread image segmentation algorithm using a convolution-based deep learning model to precisely locate and classify known anomalies. The scope of anomalies this model can classify in LPBF are recoater hopping, recoater streaking, incomplete spreading, swelling, debris, spatter, soot, super-elevation, part damage, and misprinting. The model was trained in a supervised manner using post-spread images labeled for anomalies pixel-by-pixel by field experts.

Baumgartl et al. [5] developed an image segmentation approach to identify spatter and delamination anomalies in infrared process monitoring imagery. The approach uses a simple convolution-based deep learning model with a dropout layer in-place of a fully connected classifier. The model solves a 3-class characterization problem by labeling the input as either normal, spatter, or delamination. The overall accuracy of the model was 98.9%. Additionally, gradient-weighted class activation mapping was used to visually demonstrate where class predictions were strongest in input image.

While many of these approaches demonstrate an important contribution to process monitoring in AM, they all require labeled data to train their models. In our case, we do not have labels nor know every phenomenon that can be captured in our sensor modalities. We

overcome this limitation using unsupervised approaches which reveal inter-data relationships without the need for labels or limitation in scope.

Part quality can be evaluated in different ways, but one basic method is taking porosity measurements to assess the part's density and potential breaking points. Porosity is measured using X-ray scans or cross section polishing to reveal the presence of small holes and cracks. To minimize the porosity of a part, it needs to be printed using processing parameters within an acceptable window for its specific geometry and material.

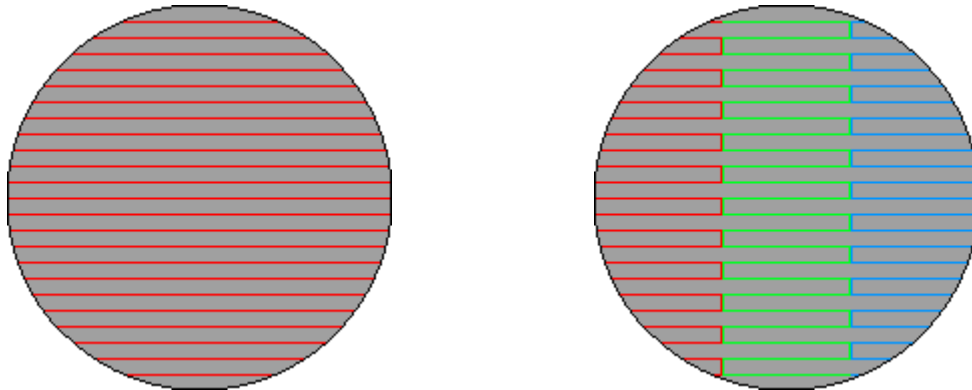
According to Kurzynowski et al. [15], there are 5 primary adjustable processing parameters that can be experimented with to optimally fabricate a part. There are powder layer thickness (L), laser power (P), laser speed (V), laser spot size (f), and hatch spacing (H). These parameters are chosen to control the Volume Energy Density (VED_H) [23] whose formula is described in Equation 2.1. It generally describes the average energy delivered within the volume of material melted by the laser spot.

$$VED_H = \frac{P}{VHL} \quad (2.1)$$

Powder layer thickness controls how many layers will be needed. Time can be saved by using fewer thicker layers, but this results in rougher surfaces and demands a higher laser power. This value ranges from 20 to 200 microns depending on material and part geometry. Laser power varies how much energy is delivered on the laser spot. Higher powers are more appropriate for thicker layers where more energy is required to melt larger volumes of material. This value usually ranges from 50 to 500 Watts (W). Laser speed must be chosen to accommodate the conductivity, viscosity, and surface tensor of the material. Faster speeds reduce the VED_H . This value usually ranges from 200 to 2000 mm/s. Laser spot size influences the laser intensity and weld spot area. Increasing the laser spot size can produce wider weld beads, however if the laser power is fixed for even larger sizes, the bead width will shrink due to lower average laser intensity. This value usually ranges from 35 to 200 microns. Hatch spacing measures the distance weld lines are separated by. The

spacing is directly proportional to the weld bead width as increasing the spacing requires a wider weld bead to maintain the necessary small overlap between welds. This value usually ranges from 10 to 200 microns.

The hatching strategy describes the laser scan pattern for a single layer. This can be very unique, but a typical strategy for bulk regions is to melt overlapping stripes and rotate their orientation by a fixed number of degrees each layer. Figure 2.2 shows the difference between continuous and striped hatches. Typically stripes are rotated with a hatch angle of 67 degrees because it maximizes the number of layers between repeated angles while creating a significant change between successive layers that prevents residual stress and material buildup in a single direction. Sometimes a hatch angle of 90 degrees is used to simplify scientific experiments.



(a) Continuous hatching. The weld lines stretch across the entire part creating inconsistent lengths as the geometry varies.

(b) Striped hatching. The weld lines are limited in length for consistency across varying geometries.

Figure 2.2: Illustration of continuous hatching vs striped hatching.

Manufacturers will use different laser surface parameters to optimize the melting of each surface of a part. These 5 parameter sets are known as in-fills, contours, up-skins, down-skins, and supports. In-fills melt bulk geometries fast and hot to reduce build time while maximizing part volume density. Contours melt layer boundaries slower to produce a sharper split from unfused powder and a smoother outer surface. Up-skins smooth-finish upward-facing surfaces with a slow hot laser to produce a glossy appearance. Down-skins

build overhangs using less energy to prevent over-heating previous layers and avoid clumping to the downward-facing surface. Supports provide a weak mechanical support and heat sink to overhangs that is easy to remove in post-processing. Figure 2.3 shows where each parameter set would be used for a part. In this thesis, we analyze coupons printed using only in-fill parameter sets.




Figure 2.3: Illustration of where laser surface parameters are used.

2.2 Image Clustering

Image clustering typically uses the same algorithms as standard clustering, but it is important to tailor the approach towards images. A naïve approach would be to cluster images in pixel space where each pixel is treated as its own independent dimension. This approach has several flaws. Spatial relationships across multiple pixels are not preserved. It is sensitive to spatial and rotational variation. The dimensionality of the pixel space is often too large making it time consuming to perform clustering. Image feature embedding solves these issues.

Image encoders can be utilized to extract higher-order features spanning several pixels and embed the presence of those features across individual dimensions. This yields a richer, compact feature space capturing high-level information that correlates better to the semantics of an image while drastically reducing the dimensionality. We use feature embedding to boost the performance of our clustering.

Recent works have been combining the clustering and feature extraction processes to improve the overall clustering [30]. We originally worked on a similar solution using a Variational Autoencoder (VAE), however due to some training complications with reconstructions, we found clustering after feature extraction was satisfactory.

2.3 Texture Analysis Using Image Preprocessing

Texture classification is an area of research under visual texture analysis. The task is to classify images of surfaces containing textural components. Textures contain varieties of wide gradients, edges, and high-frequency noise. They can vary in levels of spatial homogeneity which measures the repeatability of the texture. This means spatially in-homogeneous textures do not contain repeating patterns [31]. This is important to acknowledge because most of the LPBF in-situ monitoring data is spatially in-homogeneous thus we don't need to incorporate self-similarity techniques.

One of the early approaches to texture analysis is feature engineering of texture images using gray-tone spatial dependency matrices [12]. There was success using the piece-wise linear discriminant function method and min-max decision rule as approaches to texture classification using these texture features, however the efficiency of these features heavily rely on parameterization on the gray-tone spatial dependency matrix.

Varma et al. [28] developed a texture classification approach that utilizes 2 learning stages before the classification stage. In the first learning stage, texon dictionaries are built for each class using filter responses obtained by convolving images with filter banks. The sec-

ond learning stage generates multiple models per class by obtaining filter responses from training images and labeling those responses with the nearest texton. This generates texton frequency histograms with heights corresponding to the number of filter responses labeled with a particular texton. The classification stage computes the texton frequency histogram of the test image and classifies it with the label of the nearest neighboring model histogram. This approach using filter banks worked well and was not far off from the next step, but with the recent success of convolutional neural networks (CNN) [22], filter banks and their filter response distributions can be learned to fit many image recognition tasks.

Much research [4][8] has been using deep representation learning to perform feature extraction and embedding for texture classification. There has been some success by tailoring CNNs specifically for texture images.

Andrearczyk et al. [4] developed TextureCNN which is a modified standard CNN such that the final convolution layer is pooled instead of flattened before the fully connected layers. This discards unnecessary global spatial information while speeding up computation. In addition, all pooling layers replace max with average pool operations as [9] showed averaging is better suited for capturing texture features.

Fujieda et al. [8] demonstrated standard CNNs operate on low frequency data making them less efficient on textures. To solve this problem, they developed a WaveletCNN that uses several wavelet [27] transforms embedded into its architecture as layers to capture multi-resolution high-frequency information thus boosting performance on image classification tasks.

We use these two models as image encoders in our experiments.

Experiments

In each experiment, we use deep learning to provide us with a better understanding of process phenomena captured in the LPBF data by analyzing clusters within an embedded texture feature space. Learning to encode raw pixel images into texture representations will boost the clusterability among samples which can lead to semantically similar phenomena being closer in feature space. This is desirable because we want to understand what LPBF process phenomena are being captured by each sensor and whether our approach can effectively isolate them by clustering.

3.1 Datasets

This section describes the datasets we used and their corresponding preprocessing. The texture datasets provide a necessary foundation for our encoders to learn a compact texture representation. The LPBF dataset is our target for analysis. It contains layer-by-layer images from several sensors for a single build using a wide range of processing parameters to produce a variety of sensor responses.

3.1.1 Texture Datasets

To obtain texture representations, we leverage deep learning to extract textural features from images. For a model to learn textural features, a diverse set of texture data is required. For this reason, we considered two publicly available texture datasets, namely Kylberg and ALOT.

The Kylberg [16] dataset is a collection of texture images. It includes 28 classes with 160 unique samples per class. All samples are 576 x 576 monochrome 8-bit depth images standardized with a fixed mean of 127 and standard deviation of 40. Figure 3.1 shows a sample of each class. Notice how some textures are more spatially homogeneous like ceiling2 and scarf1 versus stone1 and wall1. Classes like lentils1 and cushion1 are made up of circular textural components while grass1 and rug1 utilize lines. High, medium, and low frequency features can be observed in floor2, ceiling1, and scarf2, respectively. We reserve this data as a baseline training set due to its smaller size and adequate textural diversity.

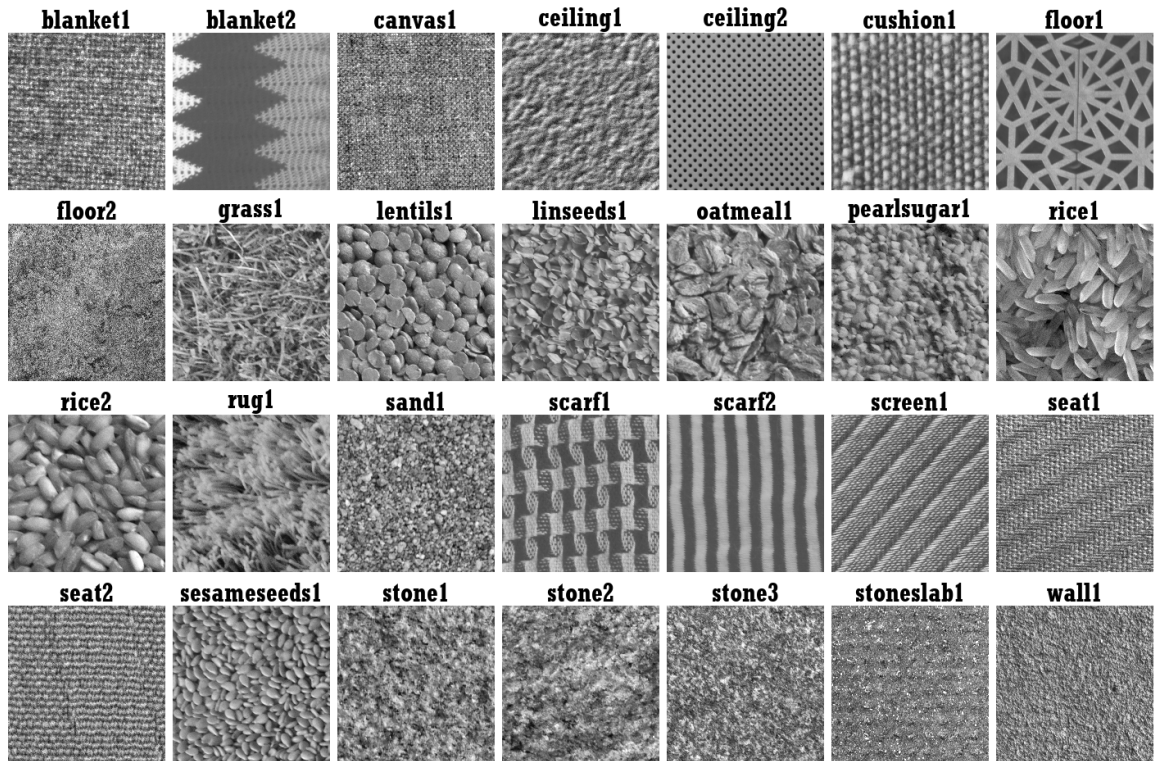


Figure 3.1: Kylberg samples from each class.

As for the preprocessing of Kylberg, we resize each sample down to 192 x 192 using Lanczos resampling and split them up into 9 non-overlapping 64 x 64 tiles. We chose these scale and split values to maximize the number of new samples while preserving an adequate resolution for detail and window area for texture variety. Figure 3.2 shows an example of this preprocessing.




Figure 3.2: Preprocessing steps for Kylberg.

The Amsterdam Library of Textures (ALOT) [6] is a huge and diverse collection of textures. It includes 250 classes with 100 lighting-augmented samples per class. There are multiple resolutions of this data, however we found using a 4X-downscaled version (384 x 256) was sufficient. Unlike the Kylberg dataset, this data is not initially normalized. Observe in Figure 3.3 the lighting across different classes plays a large role on the image contrast. It is for this reason that we perform standardization per image, and Kylberg follows suit. We reserve ALOT data for experiments that test whether the textural diversity of Kylberg is sufficient to extrapolate feature encoding to LPBF data.




Figure 3.3: ALOT samples from each class.

As for the preprocessing of ALOT, we resize each sample down to 192 x 128 using Lanczos resampling and split them up into 6 non-overlapping 64 x 64 tiles. Our reasoning is the same as Kylberg. Figure 3.4 shows an example of this preprocessing.




Figure 3.4: Preprocessing steps for ALOT.

3.1.2 LPBF In-situ Sensor Data

In this thesis we analyze LPBF in-situ process monitoring data from a single build produced by OpenAdditive. This build was made from argon-atomized Inconel-718 powder. Nine square (10 mm) coupons were built 3267 layers (135 mm) tall using 67-degree continuous hatch rotations and 243 unique processing parameter combinations varying in layer thickness, laser speed, laser power, laser spot size, and hatch spacing. Table 3.1 lists the variations in each processing parameter. These variations can produce a wide range of process phenomena that can be captured by process monitoring. In each coupon, 27 5 mm sections are assigned a random parameter set without replacement, therefore the image samples from each coupon section should have a similar texture distribution. If there is any bias in the texture distribution for individual coupons, rows, or columns in the build plane, other environment variables are suspected to be the cause.

Laser Power (W)	Laser Speed (mm/s)	Hatch Spacing (μm)	Layer Thickness (μm)	Laser Spot Size (μm)
180	400	70	30	55
230	800	90	45	86
280	1200	110	60	115

Table 3.1: Pool of processing parameters. All 243 combinations of these parameters were built by switching parameter sets every so many layers across 9 coupons.

Most approaches to process monitoring in LPBF use at least one sensor. We focus on 4 modalities provided by OpenAdditive. These are visible recoat, thermal tomography, long-wave infrared (LWIR), and spatter. Figure 3.5 shows how these sensors are arranged above the build plane. Each modality is captured by an off-axis camera, so we were given a calibration image to compute a homography matrix to calibrate camera perspective of the images into an overhead view. After calibration, each layer image is cropped to the internal area of each coupon. These 9 cropped regions are resampled up or down using Lanczos4 interpolation to 64 x 64 pixels. In the next sections, we discuss how the data for

each modality was collected and preprocessed.

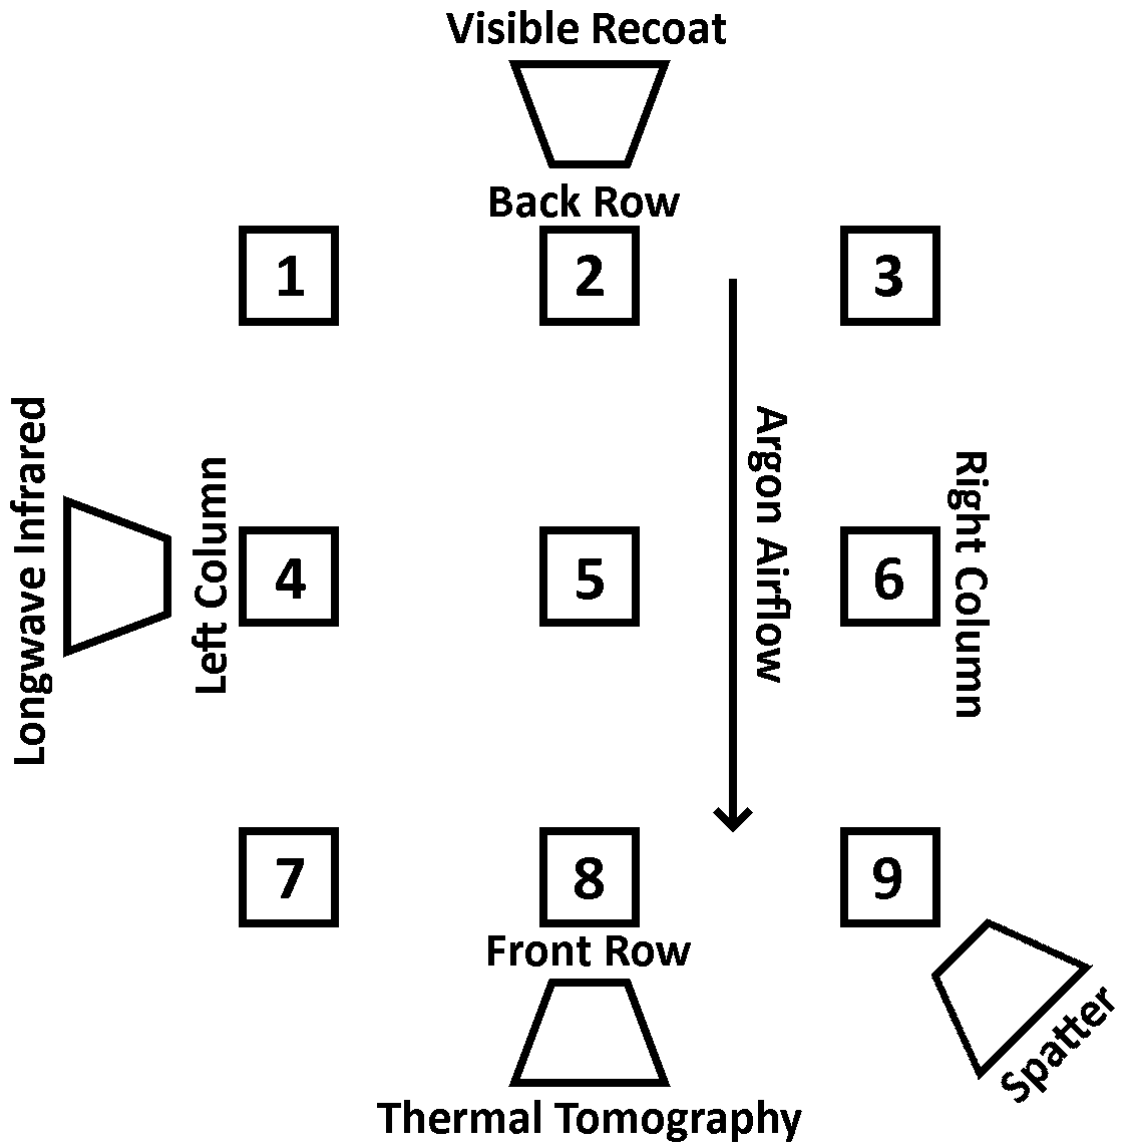


Figure 3.5: Arrangement of cameras above and coupons within the build plane.

Visible Recoat

The recoat data is captured by a Basler acA4024-29um camera with a Sony IMX226 CMOS sensor that takes 2456 x 2052 resolution monochrome pictures of the build area after spreading a powder layer and after melting it. The camera had a 3 rows and 2 columns

of dead pixels due to prior laser damage. We filled in 2 rows of dead pixels obstructing the coupon data with their nearest neighboring pixels. Figure 3.6 shows an example of powder spread and melt recoat images. This is one of the more common sensor modalities used in production. The post-spread images can reveal the uniformity of the powder spread and whether material from previous layer is protruding into the next layer. The post-melt images may reveal how defects in the powder spread affect the melted material. They also document each layer such that defects in other process monitoring data may be cross-validated with this data.




Figure 3.6: Examples of in-situ recoat images. The angled horizontal black line in both images is residue after preprocessing and laser damage to the sensor.

Thermal Tomography

The thermal tomography is captured by a Basler acA4024-29um camera with a Sony IMX226 CMOS sensor equipped with near-infrared (750-1000 nm) and neutral density filtering lenses. It is configured for long-exposure (250 ms) and low-frame-rate (4 Hz) at a 2000 x 2500 resolution. The sensor had several noisy pixels which required the images to

be cleaned using a median window convolution around the hot pixels. Figure 3.7 shows a preprocessed example layer of this data. Each coupon has a unique processing parameter set that influences the amount of energy in the welding process. This affects the emission of photons within the camera's operating wavelength spectrum. Sometimes the sensor is not calibrated to handle the full range of emissions which coincidentally washes out texture detail when the sensor response is maxed out by too many emissions. The texture in this sensor modality can be very informative because several phenomena are known to show up in thermal tomography. These include white comets (spatter), dark spots (missing powder), dark lines (recoater streaking or hopping), bright edges (delamination).




Figure 3.7: Example of thermal tomography for a single layer. The center left coupon is experiencing washout.

LWIR

The LWIR data is captured by an Optris PI 640i long-wave infrared (8-14 μm) camera. It operates at a 32 Hz frame-rate capturing temperatures between 150 - 900 C° in a 640 x 480 16-bit matrix. We convert the raw data into 8-bit images such that the pixel values 0 - 255 correspond to temperatures 150 - 900 C° while maintaining a precision error below 0.4% or about 3.5 C° . These images resemble frames capturing the heating and cooling

rates at each spatial point over the printing duration of the layer. We compile these frames using pixel-wise max to capture the peak temperatures across the build. Figure 3.8 shows a preprocessed example layer of this data. When LWIR is compiled in this way, its texture can reveal uneven distributions of the max temperature potential; however scan lines in the coupon heat maps may skew analysis to some extent. These lines are induced by passes of the laser producing hotter pixels and hatch spaces where the laser does not pass through producing cooler pixels.




Figure 3.8: Example of peak temperatures from LWIR for a single layer. The hatch angle can be observed by the lines in the coupons.

Spatter

Spatter data is captured by a Basler acA1920-155um camera with a Sony IMX174 CMOS sensor equipped with a near-infrared lens. It is configured for short-exposure ($500\mu\text{s}$) and high-frame-rate (150 Hz) at a 1100 x 1200 resolution. Since the exposure of the collection is very low, images capture the melt pool and spatter. These images are compiled into one final image for the entire layer by pixel-wise addition followed by an exponential activation function with a horizontal asymptote that bounds the sum beneath the 8-bit pixel depth (255). The specific function constants were chosen such that a fully bright pixel represented in 200+ images maps to a fully bright pixel in the compiled image. The reasoning is driven by the high frequency capture rate of images and the overlap of melt pools across images. This scaling helps reduce over-saturation from occurring thus preserving textural information. Figure 3.9 shows a preprocessed example layer of this data. Spatter compiled in this way reveals a lot of textural information. It may inform us where melt pools were larger or smaller, and where more or less spatter was ejected. There is a possibility that this data may skew analysis due to lattice artifacts produced by the hatching strategy and frame rate.




Figure 3.9: Example of compiled spatter for a single layer. Notice the coupon tails. These are hot emissions and soot being pulled away by the airflow.

3.2 Models

Our goal is to train an encoder to extract textural features from images and embed those features compactly into a vector. We leverage deep learning by training a convolutional neural network (CNN) to classify texture imagery. We implement two modified published architectures, namely TextureCNN in Section 3.2.1 and WaveletCNN in Section 3.2.2.

Across both models, we use the same fully connected layers. The first two layers use are activated with leaky ReLU (0.2 slope), the third with ReLU, and the last with log softmax. These layers are designed such that the raw feature map vector obtained from the convolutional encoder is projected into high-dimensional space, compressed through a bottle neck layer, and finally mapped to a class prediction vector which varies in dimension based on the number of classes in the texture dataset. The bottle neck layer is low-dimensional to mimic the bottle neck in an auto encoder. We refer to the activation of the bottle neck as the latent vector for the input image. This vector has 16 dimensions realizing a 2048:1 dimensionality compression ratio. We cluster LPBF imagery using latent vectors.

Both models share some properties in their convolutional encoders. Convolution layers will double the number of feature maps while preserving their height and width. These layers are activated with a 0.2-sloped leaky ReLU and average pooled using a 2 x 2 kernel. Exceptions are listed individually. All models are trained using the Adam optimizer. The learning rate is exponentially decreased from 1e-4 by 1% per epoch. While training, individual images are standardized to have a zero mean and unit standard deviation. We chose to standardize in this way because lighting can vary contrast and energy can vary average pixel intensity in LPBF sensor imagery. The rest of the augmentation includes random flips, 90° rotations, and small Gaussian noise. We stop training when losses flattens out on a log scale or overfitting starts to occur.

3.2.1 TextureCNN

We adopt the use of energy layers from TextureCNN [29] to craft a similar architecture. Our encoder uses 4 convolution layers and 4 energy layers. These energy layers are applied after each convolution layer to shortcut the filter responses from both low and high order feature maps directly to the classifier’s fully connected layers. This design is intended to boost the representation of low-level features into the latent vector. The last convolution keeps the model smaller by does not doubling the number of feature maps. Our specific

model architecture is presented in Figure 3.10.




Figure 3.10: Our TextureCNN architecture.

3.2.2 WaveletCNN

We adopt and modify the WaveletCNN [8] architecture which utilizes wavelet layers that perform multi-resolution spectral analysis to extract stronger textural features from images for texture classification. We design our encoder to utilize 4 wavelet layers and 4 convolution layers. An input image is wavelet transformed and its 1st order wavelet maps are fed as feature maps to the first convolution layer. While this is occurring, the lowest frequency wavelet map from the first wavelet layer is ran through a second wavelet layer to perform multi-resolution analysis. These 2nd order wavelet maps are appended onto the feature maps produced by the first pooling layer. This process repeats until the wavelet maps and feature maps are 4 x 4, which then we apply one last convolution layer, a 4 x 4 kernel average pooling, and flattening. Figure 3.11 shows our modified WCNN architecture.

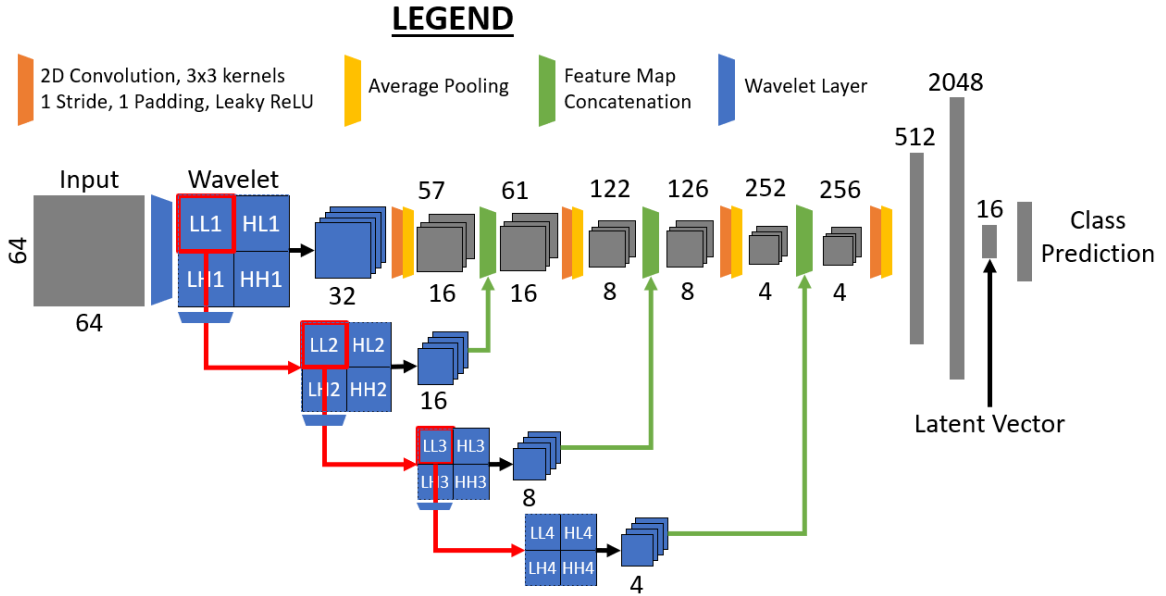


Figure 3.11: Our WaveletCNN architecture.

3.3 Fast Density Clustering

Clustering is the method we use to analyze our unlabeled LPBF data. We expect to find different process phenomena in each sensor modality manifested by a unique textural representation. By clustering these texture representations, we can discover what textures group together, understand how they correspond to process phenomena, and save time in labeling towards the development of a nearest-neighbor process anomaly classifier.

We do not know the distribution of textures in the LPBF data, so we select a multiscale clustering algorithm that does not require the number of clusters to be known beforehand. Alexandreday, et al. [3] developed a Github repo for a Fast Density Clustering (FDC) algorithm founded on kernel density maps and density graphs. Its speed comes from quick nearest neighbor searches using kernel-density trees. It supports clusters varying in size, density, and convexity. It benchmarks well with many density-based clustering algorithms and datasets in Scikit-learn library [24]. It has a few parameters but is reasonably robust, so we only consider tuning the noise threshold which merges overlapping clusters. FDC is

most efficient on low-dimensional datasets. Since we learn texture representations in only 16 dimensions, we find this algorithm fills our needs.

3.4 UMAP Dimensionality Reduction

Uniform Manifold Approximation and Projection (UMAP) [18] is a dimensionality reduction algorithm. Points far apart in the data space will also be far apart in the mapped space. It is similar to t-Distributed Stochastic Embedding (t-SNE) but is faster, consumes less memory, and preserves more global structure. We utilize the UMAP Github repo [21] to project our dataset feature spaces onto a 2D plane for scatter plot visualization. This helps us grasp how texture feature spaces cluster for a given dataset. While UMAP can be used as a preprocessing step for clustering [19], we do not find it necessary since our texture representation is already low-dimensional.

It is critical to note that the design of UMAP, like tSNE, has a few limitations that impact interpretability. UMAP is a stochastic algorithm and requires some parameter tuning. Leland et al.[18] describe two primary parameters; the minimum distance controls how close points are packed together while the number of neighbors controls how much global structure is preserved. The broader topology of data can be preserved by focusing more on global structure from increasing the number of neighbors and increasing the minimum distance to expand neighboring clusters until they connect. Additionally, Coenen et al. [7] showed for several toy datasets that cluster sizes and exact distances between clusters usually mean nothing and cannot be used to draw conclusions while comparing clusters in a visualization. We acknowledge these findings and use caution when interpreting our results.

We illustrate the utility of UMAP visualization on MNIST [17]. MNIST contains 70,000 hand-written digits encoded as 28 x 28 grayscale images. By flattening these images into 784-pixel vectors, we can visualize how the data groups together using UMAP to project

the pixel space onto a 2D scatter plot as shown in Figure 3.12. The clustering is considerably good despite no prior feature embedding. This performance can be accredited towards the simplicity of MNIST. Samples do not differ much in translation nor rotation, which explains why semantically equal samples will be close in pixel space.

There's a few observations to point out here. Clusters 4, 7, and 9 are very close because their digits share a right vertical line. This also occurs for clusters 3, 5, and 8 because their digits share a lower curve. Clusters 0 and 1 may not necessarily be far apart because their pixel representation is vastly different, but because the 2D space cannot preserve topological properties, such as wrapping, that show up in higher-dimensional spaces.




Figure 3.12: UMAP Visualization of MNIST flattened image vectors parameterized with a minimum distance of 0.1 and 15 nearest neighbors. The points are colored by their class label. The homogeneity of color in each cluster indicates most samples in tight proximity share the same semantic label which we view as a desirable clustering.

Results and Discussion

In this chapter we analyze each LPBF sensor modality in 4 experiments. In each experiment we choose a model and a texture dataset to train that model. For Kylberg, we use a 9:1 train-test split that yields 1296 samples per class in the training set. We use a 5:1 train-test split for ALOT which yields 500 samples per class in the training set. We train a model on one of these datasets until the loss plot flattens or shows signs of overfitting. Afterwards, we run each LPBF sensor modality data through the model to obtain feature vectors. We perform FDC on these feature spaces, examine the images within each cluster, and visualize these clusters on a scatterplot using UMAP.

Table 4.1 lists the experiment classification scores. Our TCNN model for Kylberg performs similarly in [4] despite our preprocessing steps and bottleneck layer. Our WCNN does not perform as well on Kylberg, however this could be a result of early stopping. The ALOT models have noticeably lower scores. This is most likely explained by the fact that ALOT has nearly 9x more classes than Kylberg.

Dataset	Model	Accuracy	Epochs
Kylberg	TCNN	0.9797	85
Kylberg	WCNN	0.9524	50
ALOT	TCNN	0.7306	67
ALOT	WCNN	0.7428	88

Table 4.1: Classification accuracies of each experiment.

These scores are reasonably important, however we just use them to gauge how well the encoders are learning. We are primarily interested in using these encoders to extract

and embed textural features from new data which we can cluster and visualize on a 2D scatter plot projected by UMAP.

The main discussions here highlight clustering the feature spaces of each dataset. We start by analyzing the texture datasets in Section [4.1](#) to assess the utility of our approach. Afterwards we apply our approach on LPBF data in Section [4.2](#).

4.1 Texture Datasets

In this section we analyze the feature spaces of each texture dataset. Given a texture dataset, we train a TCNN and WCNN on that dataset. After training each model, we transform the texture set images into feature vectors using those models to obtain two feature spaces. These spaces are clustered using FDC and projected to a plane using UMAP for scatter plot visualization. We analyze Kylberg in Section [4.1.1](#) and ALOT in Section [4.1.2](#).

4.1.1 Kylberg

In this section we cluster the Kylberg feature space from a TCNN and WCNN trained on Kylberg. We begin our analysis by plotting the average activation of each feature dimension for both models in Figure 4.1. We observe not all the dimensions are being used for both models. The Kylberg-TCNN has 5 dead dimensions (dimensions 7 and 15 are nonzero), while the Kylberg-WCNN has 3. This behavior might be explained by the ReLU activation on the bottleneck in both models. Nevertheless, both models were still able to perform well even with 11 and 13 dimensions. We note that some dimensions have different ranges than others. Whenever we cluster data using the Kylberg-TCNN or Kylberg-WCNN texture representation, we drop the dead dimensions and normalize the rest. We make the assumption all dimensions are of equal importance because it is beneficial for clustering.

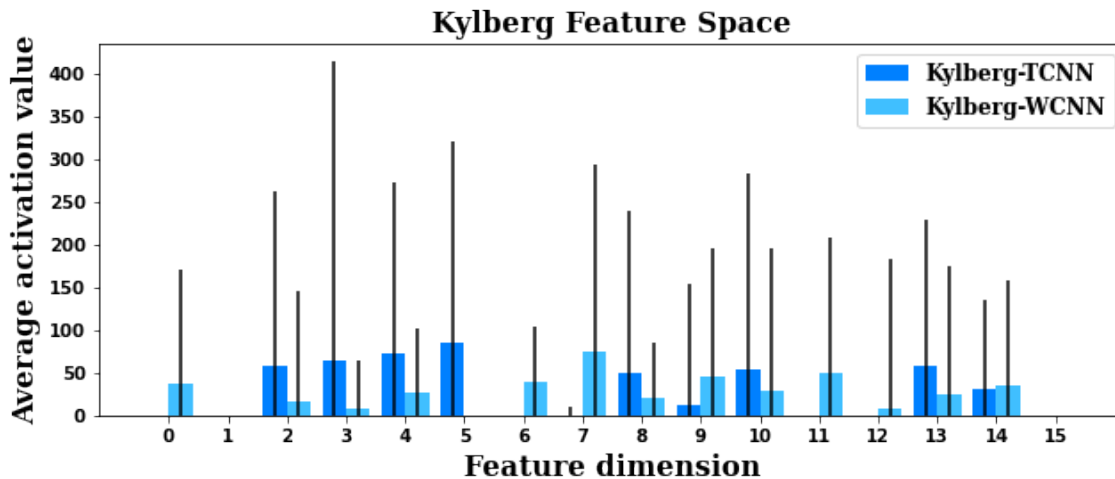


Figure 4.1: Plot of the average activation in each feature dimension from Kylberg for both models trained on Kylberg. Lines indicate the range of the activations.

With 28 classes, we would expect to see 28 clusters assuming the texture representations of each class are distinct. For both models, we show the number of clusters Kylberg breaks into while varying the noise threshold in Figure 4.2.

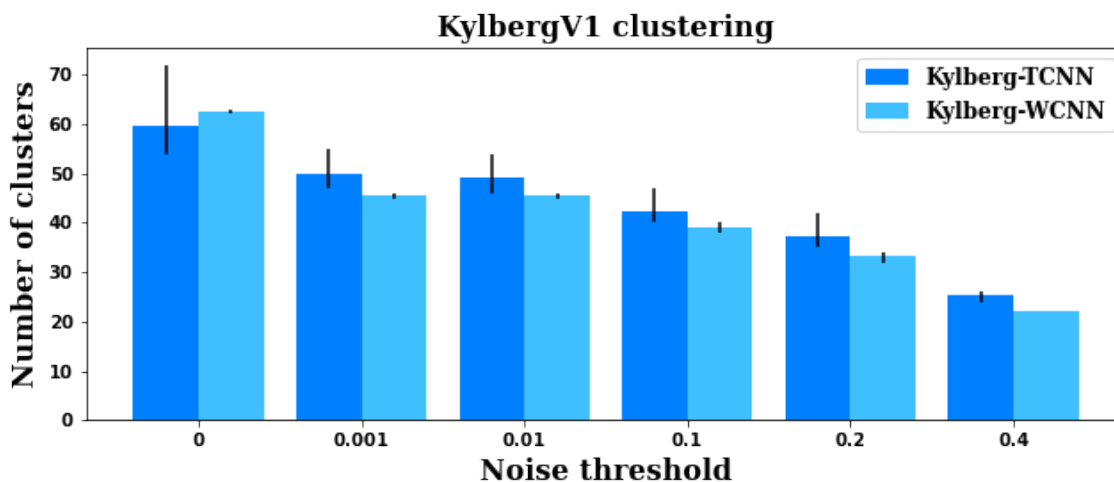


Figure 4.2: Plot of the average number of clusters Kylberg breaks into at different noise thresholds for two experiments. The range over 5 runs is indicated by a line on each bar.

With a noise threshold of 0, the number of clusters exceeds the number of classes two-fold. Both models produce similar number of clusters, so we cluster using a noise threshold of 0.2 for both models because the number of clusters produced is closer to the number of classes. Higher noise thresholds are not considered because they result in fewer clusters than there exist classes. In the following sections, we discuss the Kylberg-TCNN and Kylberg-WCNN experiments.

Kylberg-TCNN

Beginning with the TCNN, we cluster with a noise threshold of 0.2 to produce 35 clusters. These clusters are relabeled sorted by size and plotted in Figure 4.3.

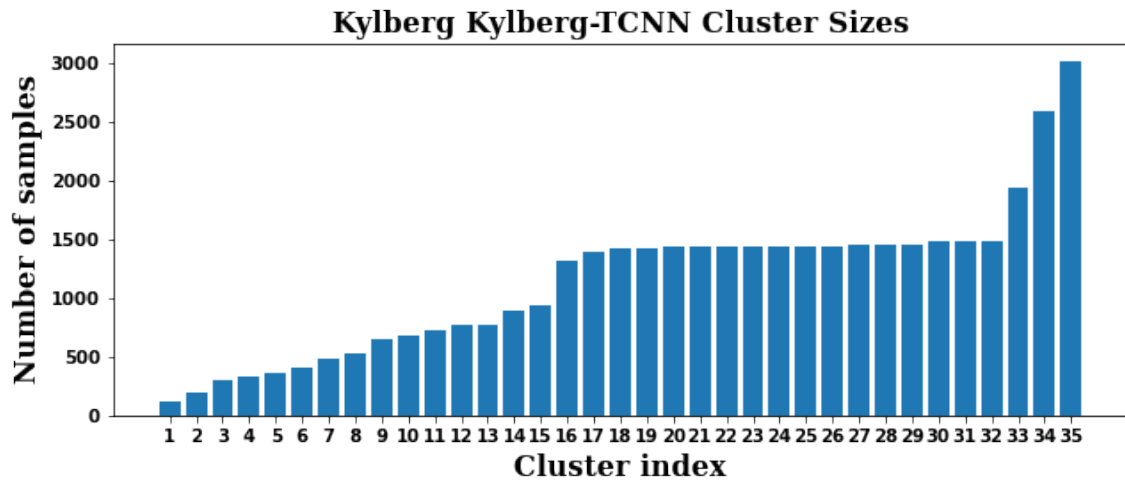



Figure 4.3: Kylberg cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.2.


From this graph we notice an imbalance in cluster sizes; however, the sizes of clusters 16 to 32 follow the expected uniform distribution. Larger clusters 33 to 35 are not as homogeneous as they must contain several samples from neighboring classes, while smaller clusters 1 to 15 likely contain intra-class variations. We explore these clusters by showing randomly sampled images from each cluster in Figure 4.4.



(1) Floor1

(2) Blanket2


(3) Scarf2



(4) Ceiling2

(5) Scarf2


(6) Ceiling2



(7) Seat2

(8) Seat1


(9) Stoneslab1



(10) Ceiling2

(11) Blanket2


(12) Stoneslab1



(13) Blanket2

(14) Seat1


(15) Seat2



(16) Floor1

(17) Grass1, Rug1


(18) Sand1



(19) Canvas1

(20) Sesameseeds1


(21) Cushion1



(22) Screen1

(23) Blanket1

(24) Lentils



(25) Rice2

(26) Wall1

(27) Stone3

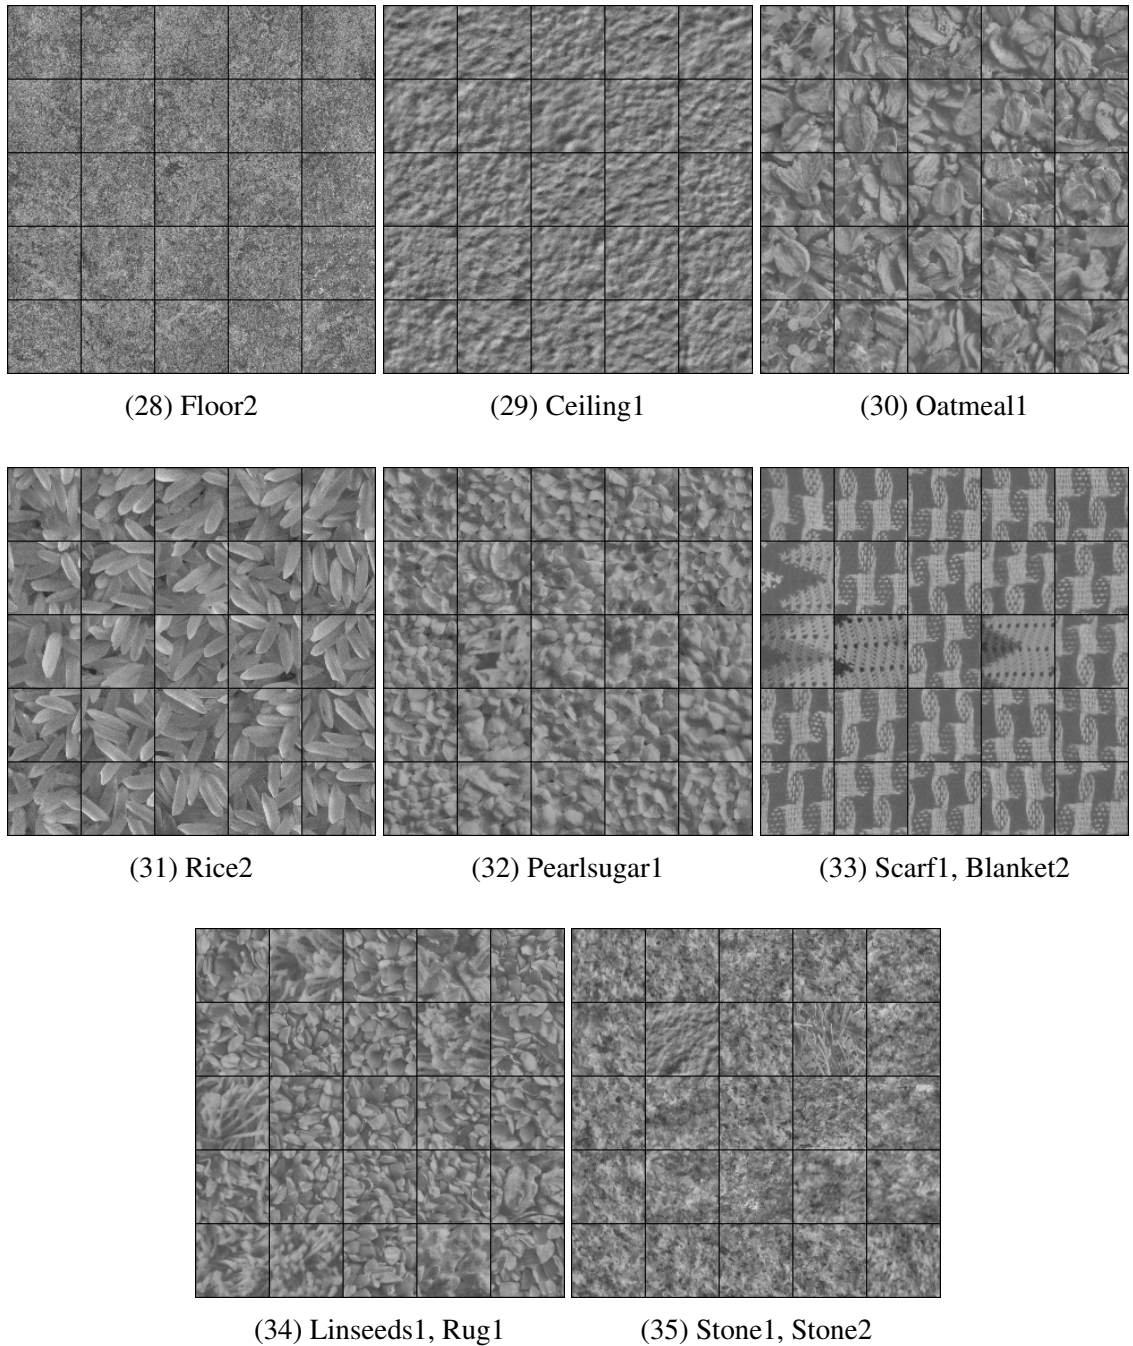


Figure 4.4: Normalized images sampled from each cluster in the Kylberg Kylberg-TCNN feature space.

We dive into the smaller clusters first. All of these correspond to classes with a strong structural pattern. Scarf2 breaks up among clusters 3, 5, and 13 based on the where the pattern was cropped. Cluster 3 features an extra light stripe, while cluster 5 features an extra

dark stripe. Cluster 13 is a mix between featuring equal number of the stripes. Ceiling2 breaks up into clusters 4, 6, and 10 based on the scale. Blanket2 splits into clusters 2, 11, and 33 due to normalization and cropping. Cluster 2 features small bright areas against a darker background, while clusters 11 and 33 feature a mix of lighter and darker zigzags. Floor1 splits into clusters 1 and 16 with the majority in the latter due to cropping. Cluster 1 captures a unique part of the pattern that is overall brighter than cluster 16 containing more dark spaces. Seat1 is spread across clusters 8 and 14 by a subtle difference. Cluster 8 has tighter, brighter stripes, while they blend in more in cluster 14. Seat2 splits into clusters 7 and 15; however we cannot confidently conclude what the differentiating feature is. Finally, Stoneslab1 splits nearly in half into clusters 9 and 12, but once again a reason is not readily obvious.

We discuss the other clusters next. Clusters 18 through 29 and 31 appear homogeneous. The other clusters show signs of noisy segmentation. Cluster 35 combines both Stone1 and Stone2 classes as they share similar peak-value distributions and frequencies. Cluster 17 shows a mixture of samples from Grass1 and Rug1. This is likely caused by the structural similarity between the rug fibers and grass blades. Cluster 30 primarily contains Oatmeal1, but some samples from Grass1 are present. This could be accredited to the similarity between clovers and oatmeal flakes. Cluster 32 is made up of Pearlsugar1 with a few samples from Rug1 and Oatmeal1 likely due to the similarity in pixel intensity distributions. Cluster 33 includes a few Blanket2 samples within a primarily Scarf1 cluster. The small holes and low-frequency dark and bright regions appear to be common features between these classes. Cluster 34 is a larger composed primarily of Linseeds1 with some Rug1, Pearlsugar1, and Oatmeal1. Cluster 35 primarily groups Stone1 and Stone2 together as they are very similar in appearance.

From clusters 17, 30, 32, 34, and 35 we notice most of the segmentation noise is rooted among their classes. To confirm this, we visualize the Kylberg feature space by applying UMAP dimensionality reduction. We varied UMAP's minimum distance and number

of neighbors parameters until we found a satisfactory projection. We show the Kylberg-TCNN visualization colored by class and cluster in Figures 4.5, 4.6, respectively




Figure 4.5: UMAP visualization of the Kylberg Kylberg-TCNN feature space colored by class. Minimum distance of 0.25 and 300 neighbors were used.




Figure 4.6: UMAP visualization of the Kylberg Kylberg-TCNN feature space colored by cluster number. Minimum distance of 0.25 and 300 neighbors were used.

From Figure 4.5 we see most classes are well-isolated within the projection with similar classes neighboring each other. Most of the spatially homogeneous classes reside on the border of the projection likely due to their unique texture representation being spatially distant from others. From Figure 4.6 we notice most density-based clusters are seman-

tically homogeneous. However, among the tight neighborhood of Linseeds1, Oatmeal1, Pearlsugar1, Grass1, and Rug1, the majority of semantic inhomogeneity comes from the Grass1 cluster extending itself into its neighboring Rug1 cluster. Additionally, Stone1 and Stone2 pack together in one cluster which partially extends into the tail of the Grass1 and Rug1 cluster. Split classes such as Seat1, Seat2, Ceiling2, Floor1, Scarf2, Stoneslab1, and Blanket2 have differently colored regions representing the clusters that split each class.

Kylberg-WCNN

Finishing with the WCNN, we cluster using a noise threshold of 0.2 yielding 29 clusters. These clusters are relabeled sorted by size and plotted in Figure 4.7.





Figure 4.7: Kylberg cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.2.

The distribution depicts a similar imbalance with 3 larger clusters. The sizes of clusters 17 to 26 follow the expected uniform distribution with clusters 27 to 29 being a bit larger. In the TCNN experiment, the biggest cluster contained nearly 3000 samples, but the largest 3 clusters in this distribution surpass that. To understand why, we explore these


clusters by showing randomly sampled images from each cluster in Figure 4.8.



(1) Blanket2

(2) Blanket2


(3) Screen1



(4) Scarf2

(5) Screen1


(6) Scarf2



(7) Ceiling2

(8) Stone3


(9) Seat1



(10) Blanket2

(11) Scarf2


(12) Seat1



(13) Screen1

(14) Wall1


(15) Stone3




(16) Ceiling2

(17) Rice2


(18) Canvas1




(19) Ceiling1




(20) Seat2




(21) Cushion1




(22) Floor1




(23) Blanket1




(24) Lentils1



(25) Sesameseeds1



(26) Rice1



(27) Scarf1

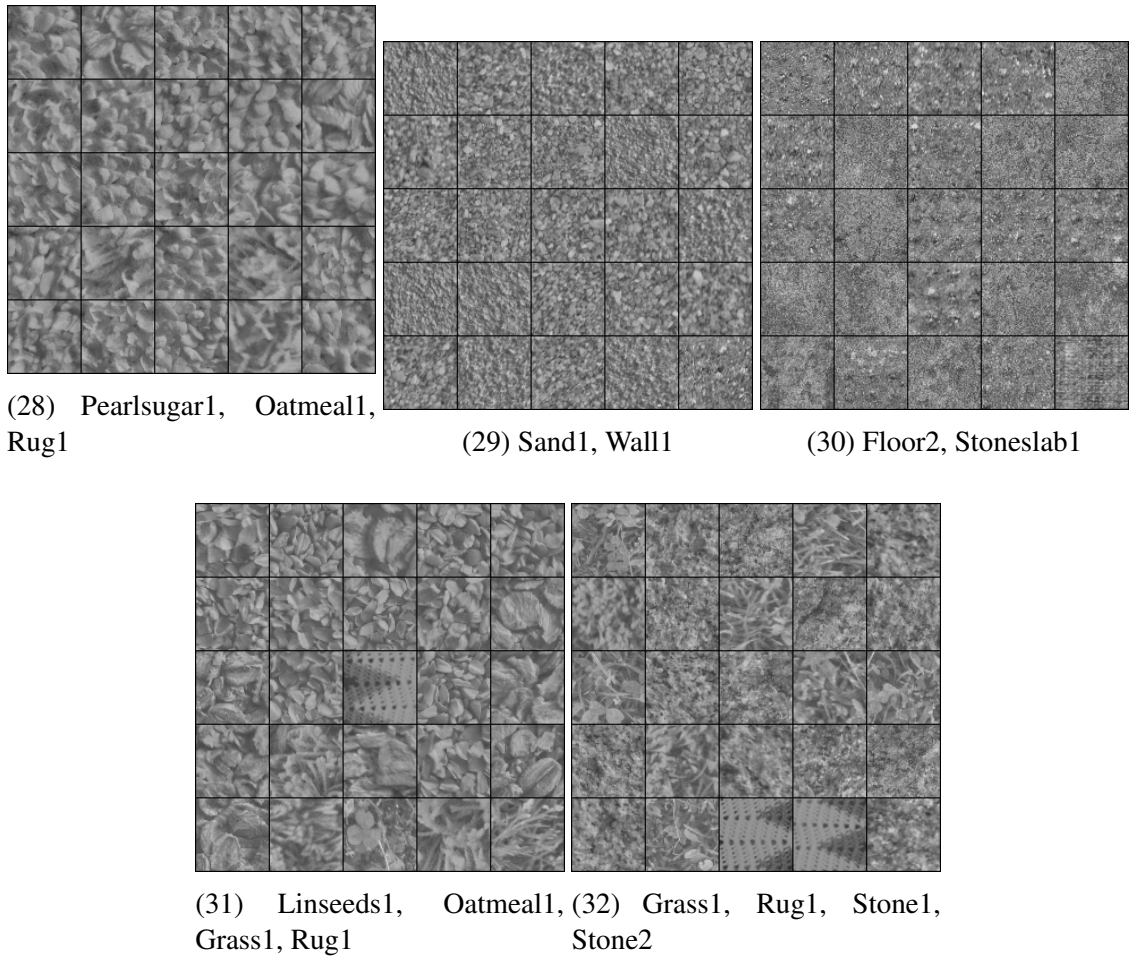


Figure 4.8: Normalized images sampled from each cluster in the Kylberg Kylberg-WCNN feature space.

We dive into the smaller clusters first. Like before, Scarf2, Ceiling2, Blanket2, Seat1 are splitting into multiple clusters; however Stoneslab1, Floor1, and Seat2 do not split. Screen1 splits across clusters 13, 5, and 3 with each respective cluster having sharper lines. Stone3 breaks up into clusters 8 and 15 being slightly rougher and smoother, respectively. Wall1 splits into clusters 14 and 29 with coarser samples being grouped in the latter cluster with Sand1.

We discuss the other clusters next. Clusters 17 through 26 appear homogeneous. The other clusters show signs of noisy segmentation. Cluster 27 primarily contains Scarf1 with some samples from Blanket2. Cluster 28 shows a similar mix to cluster 30 from the TCNN. Cluster 29 primarily contains Sand1, but also some of the coarser Wall1 samples. Cluster

30 combines both Floor2 and Stoneslab1 due to their high-frequency low-level features. Cluster 31 is similar to cluster 34 from the TCNN. Finally, cluster 32 being the largest shows the Grass1 and Rug1 cluster combined with the Stone1 and Stone2 cluster from the TCNN projection. To confirm this, we show the Kylberg-WCNN visualization colored by class and cluster in Figures 4.9, 4.10, respectively.




Figure 4.9: UMAP visualization of the Kylberg Kylberg-WCNN feature space. Minimum distance of 0.25 and 300 neighbors were used.

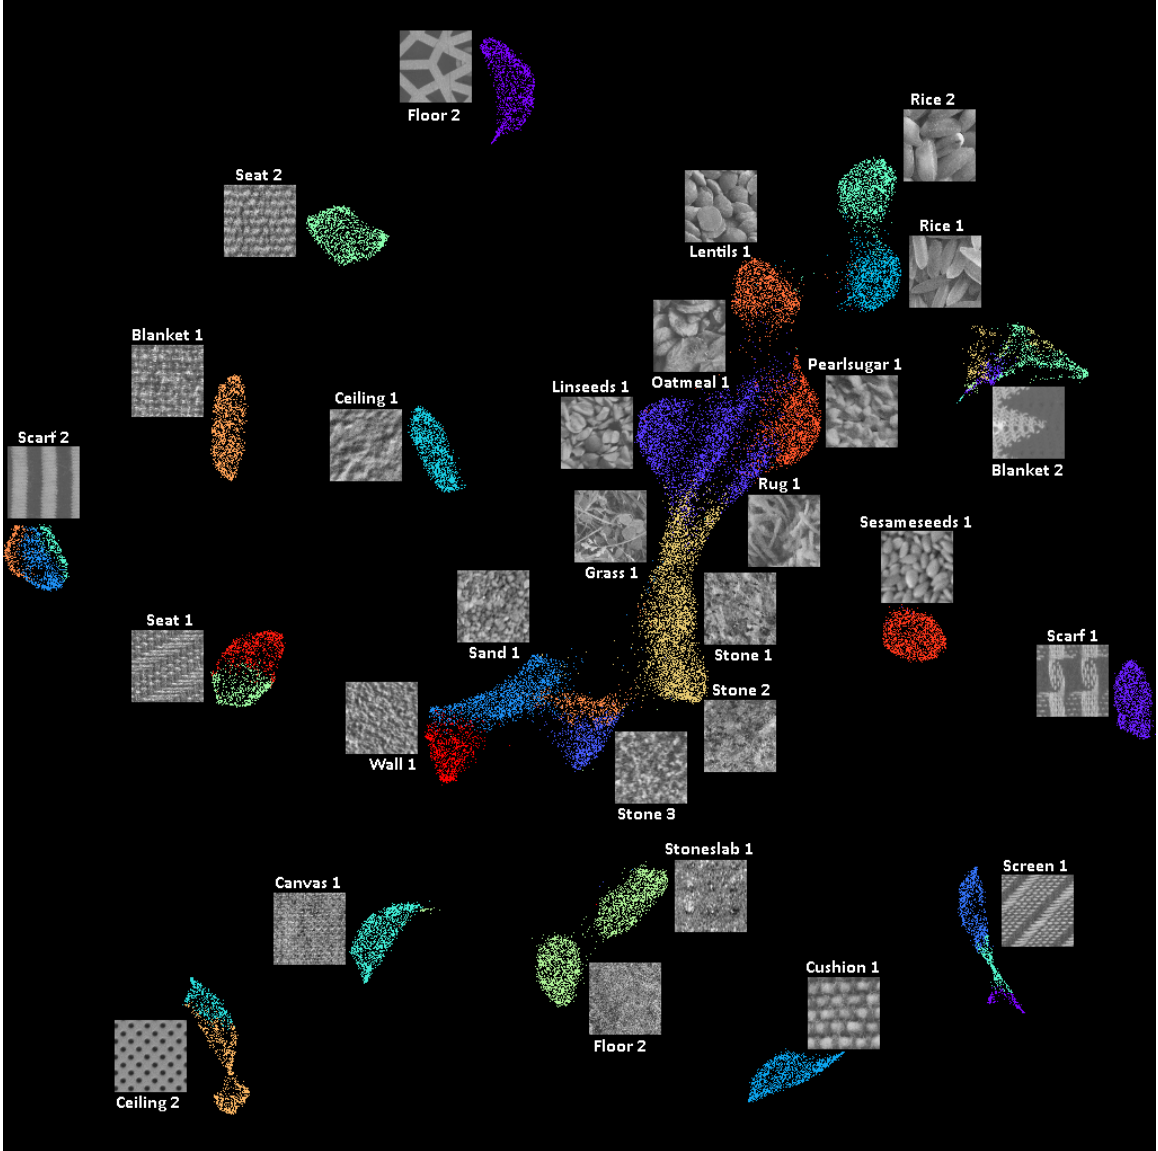


Figure 4.10: UMAP visualization of the Kylberg Kylberg-WCNN feature space. Minimum distance of 0.25 and 300 neighbors were used.

From Figure 4.9, we see a similar arrangement of clusters from Figure 4.5. The main differences are Stone1 is attached to Grass1, Sand1 sits between Stone3 and Wall1, and Oatmeal1 is stretched towards Linseeds1 and Grass1 but away from Rug1 and Pearlsugar1. From 4.10 we observe the clustering splits Scarf2, Ceiling2, Seat1, Screen1, Blanket2, and Stone3. We see Sand1 over extends into Wall1. While UMAP split Stoneslab1 and Floor2 into two neighboring clusters, they were grouped into one cluster in high-dimensional

space. This is likely caused by the noise threshold being high enough to combine the overlapping clusters. The major cluster group breaks into 3 clusters. Pearlsugar1 captures a few of its neighbors but is otherwise well isolated. Oatmeal1, Linseeds1, and half of Rug1 group together. Finally Grass1, Stone1, Stone2, and some of Rug1 cluster together. This confirms our initial observations.

Summary

In both experiments, most classes segmented cleanly in both the clustering and the projection. Similar classes and clusters were often neighbors in the projection. This was especially the case for Lentils1, Rice1, and Rice2 because their higher-order features described piles of round smooth objects. Some classes such as Stone1 and Stone2 neighbored each other so closely that neither clustering could distinguish between them. Although both experiments produced similar results, there were some differences. The TCNN clustered classes semantically better than the WCNN; however, this might be related to its higher classification score. The WCNN prioritized grouping classes based on the frequency level of their textural components. We list a few examples of this. Floor2 and Stoneslab1 both sharing high-frequency low-level noise clustered together in WCNN experiment but not in the TCNN-experiment. Sand1 acted as an interpolation between Stone3 and Wall1 because it is a gradual transition in frequency with a similar level of coarseness. This concludes our analysis of Kylberg. We discuss our results from ALOT in the next section.

4.1.2 ALOT

In this section we cluster the ALOT validation set feature space from a TCNN and WCNN trained on ALOT. We begin our analysis by plotting the average activation of each feature dimension for both models in Figure 4.11. Similar to the Kylberg models, both models have dead dimensions with varying ranges. The ALOT-TCNN representation does not use dimensions 3, 9, 13, and 15, while the ALOT-WCNN representation does not use dimensions 3, 4, 6, and 12. We drop the dead dimensions and normalize the rest.




Figure 4.11: Plot of the average activation in each feature dimension from ALOT for both models trained on ALOT. Lines indicate the range of the activations.

With 250 classes, we expect to see a fewer clusters due to the similarity between many classes. For both models, we show the number of clusters ALOT breaks into while varying the noise threshold in Figure 4.12.

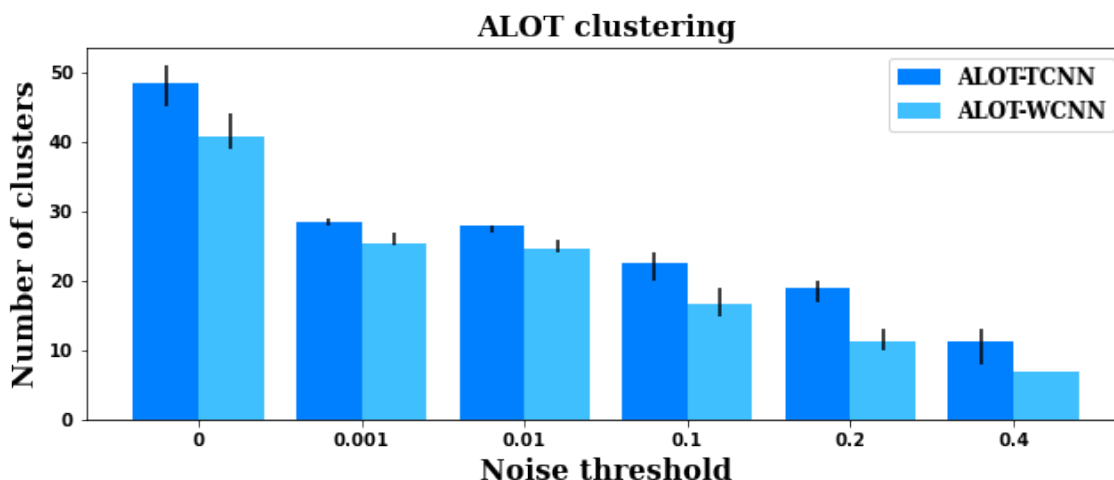


Figure 4.12: Plot of the average number of clusters ALOT breaks into at different noise thresholds for two experiments. The range over 5 runs is indicated by a line on each bar.

We observe a significantly smaller number of clusters indicating the presence of overlapping features in multiple classes. The WCNN consistency produces slightly fewer clusters than the TCNN implying the WCNN feature spaces do not break apart well enough for FDC to isolate as many distinct density neighborhoods. We later learn the frequency analysis performed by the wavelet layers is less effective at learning a well-separated embedding for textures with large shapes and solid regions. Due to the size of ALOT, we change our criteria for selecting the noise threshold. For both models, we select a noise threshold of 0.01 because the number of clusters produced is most similar between models.

We begin our analysis on the TCNN. Clustering the ALOT-TCNN feature spaces yields 28 clusters. We plot the cluster sizes in Figure 4.13.





Figure 4.13: ALOT cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.01.

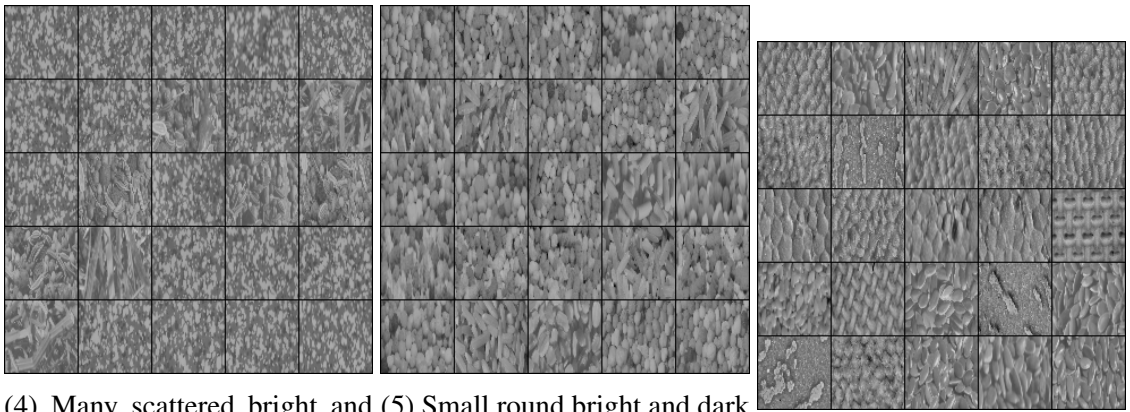
We observe nearly half of the data is thrown into one cluster, while smaller clusters 1 to 14 likely isolate a few classes. We are primarily interested in the larger clusters as they contain more samples that will span more classes that could describe a broader textual group. Figure 4.14 shows images randomly sampled from each cluster and labeled by a general textual description. We inspect these images to get a better evaluation on the textual homogeneity per cluster.



(1) Square grid

(2) Scattered tendrils


(3) Foil-like



(4) Many scattered bright and dark regions

(5) Small round bright and dark objects


(6) Small round objects




(7) Most bright, some dark objects

(8) Coarse, several bright spots


(9) Line patterns




(10) Fine bright and dark spots




(11) Round objects, dark spots




(12) Bright and dark round objects




(13) Low-frequency, round objects




(14) Large objects / shadows, no bright reflections




(15) Round objects with bright reflections




(16) Large lattices



(17) Coarse, bright and dark spots




(18) Small lattices



(19) Gears and lattices

(20) Sand-like


(21) Fine structured patterns



(22) Large round objects

(23) Rods and grooves


(24) Coarse and rough



(25) Fine and flat

(26) Edges, shadows, holes

(27) Wide variety, rough



(28) Low frequency, solid regions

Figure 4.14: Normalized images sampled from each cluster in the ALOT ALOT-TCNN feature space.

Most clusters group classes based on component shapes and sizes, spatial structural homogeneity, and distribution of bright and dark spots. Clusters 6, 9, 11, 13, 16, and 18 contain more samples with a repeating pattern. Clusters 10, 17, 20, 21, 24, and 25 contain finer samples with tiny features. Clusters 1, 2, 5, 6, 7, 9, 11, 12, 13, 14, 15, 16, 19, 22, 24, 26, 27, and 28 contain coarser samples with larger shapes, shadows, and solid regions. To see how these classes group together, we color and partially label a UMAP projection of this ALOT feature space by cluster in Figure 4.15. We labeled the larger clusters and only a few of smaller ones to keep the visualization cleaner.

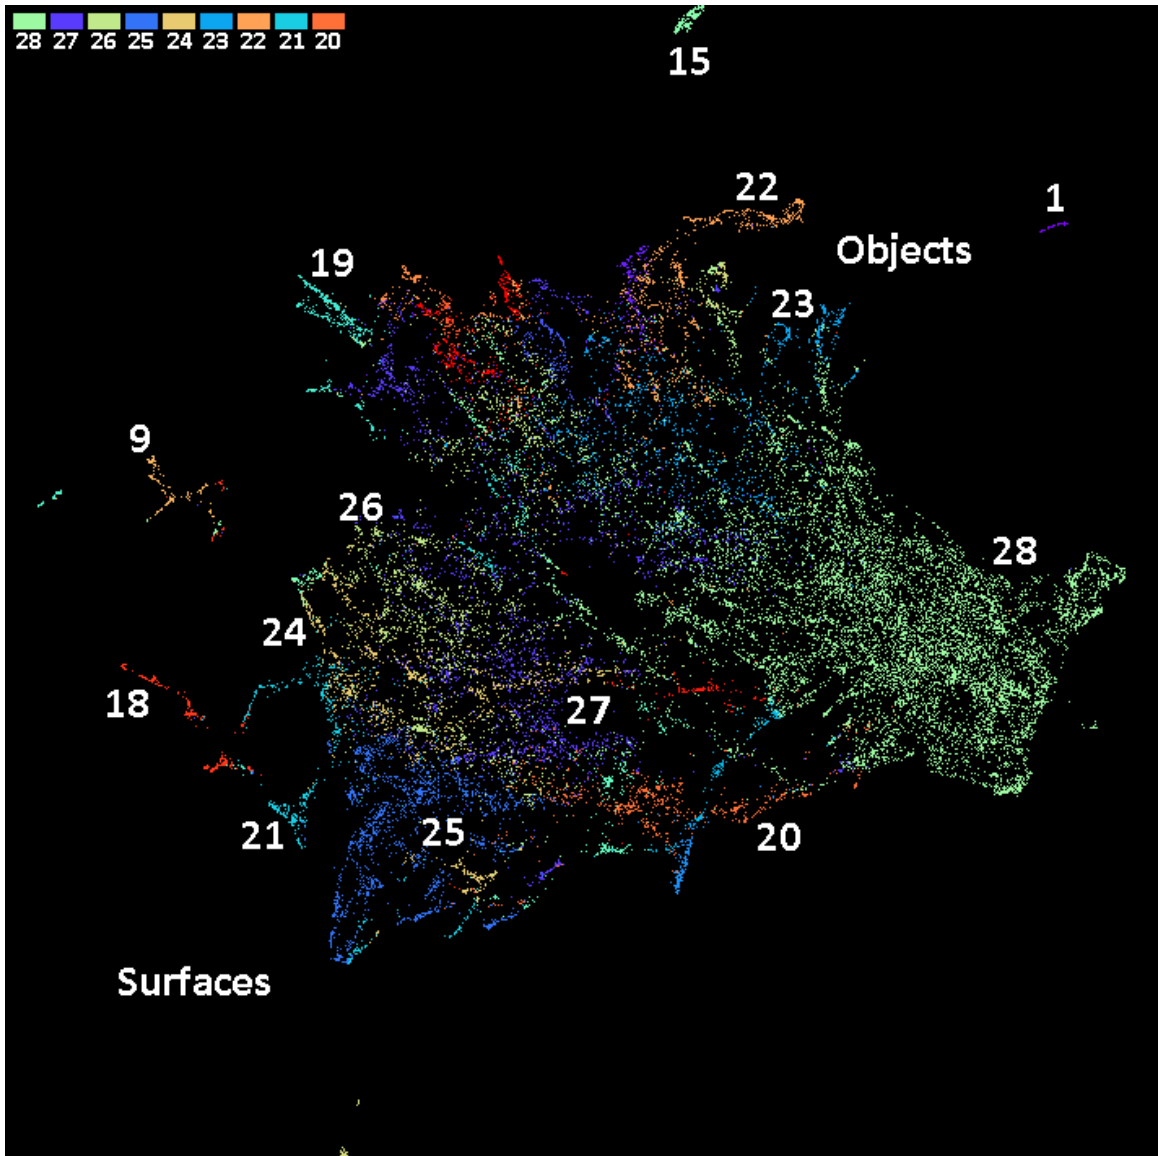


Figure 4.15: UMAP visualization of the ALOT ALOT-TCNN feature space colored by cluster. Minimum distance of 0.10 and 15 neighbors were used.

From this visualization, the classes split into two major categories. The left half is composed of surfaces with finer features, while the right half is composed of large objects with various shapes, overlap, and shadows. Each major category breaks into smaller distinct groups that vary in other features. For the surfaces, levels of roughness, noise, and spatial homogeneity are distinguishing factors. For the objects category, shape, size, and lighting are distinguishing factors.

Many clusters are contiguous in the projection, however some, like cluster 27, are scattered across the surface and object categories. We do not necessarily expect clusters in feature space to be contiguous in the projections due to the nature of dimensionality reduction and the differences between UMAP and FDC. We note most samples from cluster 27 appearing in the object category are piles of small uniformly-shaped objects, such as pellets. These piles may resemble similar rough non-spatially homogeneous surfaces in feature space. Clusters 1, 9, 15, and 18 break away from the major connected components in the projection. We suspect their unique textural representation and homogeneity may be the cause. A large portion of the objects category groups into cluster 28 where further splitting could be realized. After inspecting cluster sizes with a noise threshold of 0, we found cluster 28 only decreases in size by 20%; therefore we rule out the noise threshold as a cause for cluster 28 being so large. This leads us to believe the textural representation of the TCNN might be entangled or not segmented well enough for FDC to identify smaller overlapping dense regions in feature space.

We move onto the analysis of WCNN for comparison. Clustering the ALOT-WCNN feature space yields 24 clusters. We plot the cluster sizes in Figure 4.16.






Figure 4.16: ALOT cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.01.


We observe a non-uniform distribution similar to the previous one, except the tail does not decay as gradually. The two largest clusters likely have lower textural homogeneity as they are even larger than those observed in the TCNN experiment. Clusters 20 to 22 are sized uniformly within a size range that could capture general textural groups with higher textural homogeneity. We gauge this on the fact that each class is represented by 100 samples in the validation set which ideally could correspond to about 15 similar classes being clustered together. The rest of the clusters 1 to 19 are small with high textural homogeneity as they contain only a few classes. We inspect these clusters closer by randomly sampling and describing images from each cluster in [Figure 4.17](#).




(1) Denim




(2) Bright and dark small round objects




(3) Smooth dull round objects with shadows




(4) Tiny round objects




(5) Solid areas, round and straight edges




(6) Small dark spots




(7) Flat, scattered bright and dark areas




(8) Round glossy objects




(9) Smooth and powdery objects



(10) Shadows, smooth and powdery (11) Chunky objects with dark shadows (12) Sharp edges split bright and dark regions



(13) Linear structural patterns (14) Small elliptical objects (15) Rough surfaces with holes



(16) Round shiny objects (17) Shapeless, smooth with specks (18) High-frequency small components

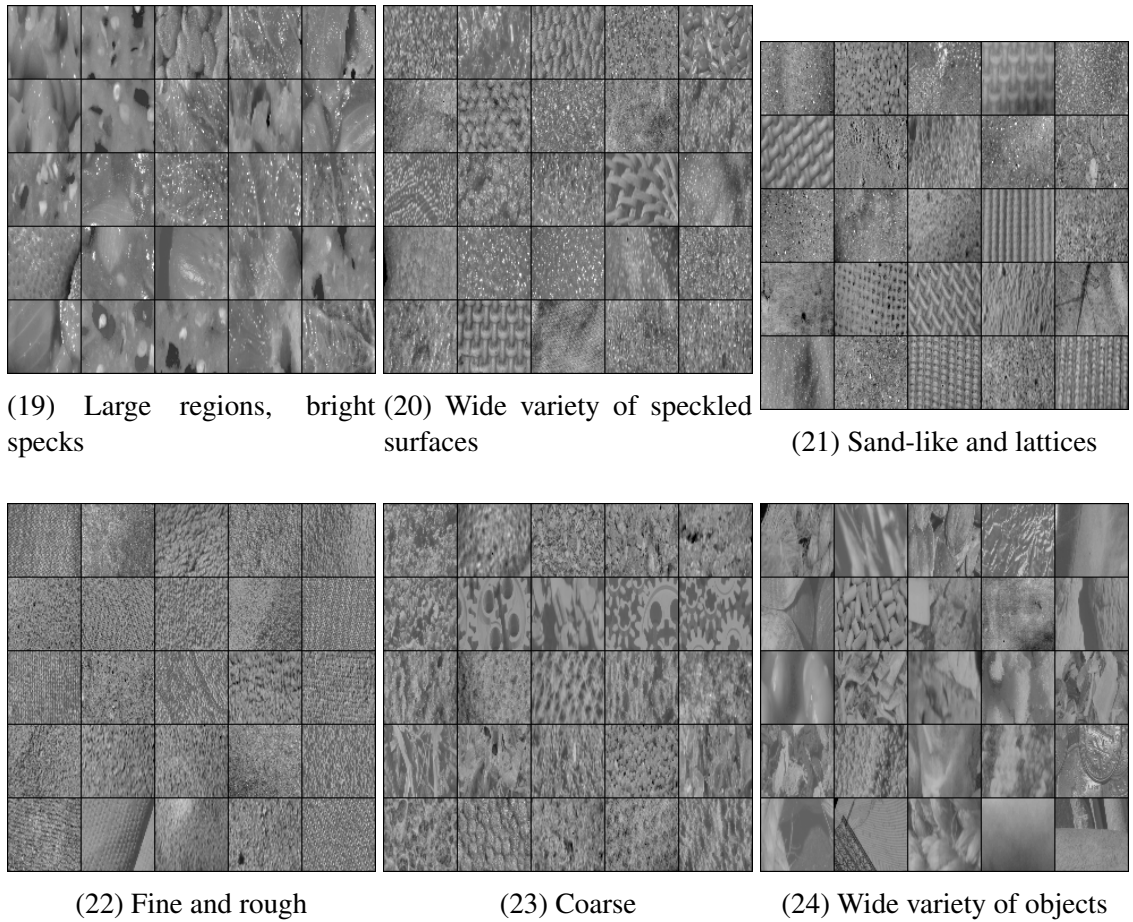


Figure 4.17: Normalized images sampled from each cluster in the ALOT ALOT-WCNN feature space.

In comparison to the TCNN experiment, we obtain more small clusters isolating samples from a couple classes. Clusters 1, 5, 6, 7, 13, 15, 18, and 20 are generally surfaces, while clusters 2, 3, 4, 8, 9, 10, 11, 12, 14, 16, 17, and 19 are objects. Clusters 21 and 22 generally describe grainy surfaces, while cluster 23 contains coarser surfaces and some objects. Nearly all object classes are located in cluster 24. Like before, we color a UMAP projection of this ALOT feature space by cluster in Figure 4.18 to demonstrate where clusters are located.

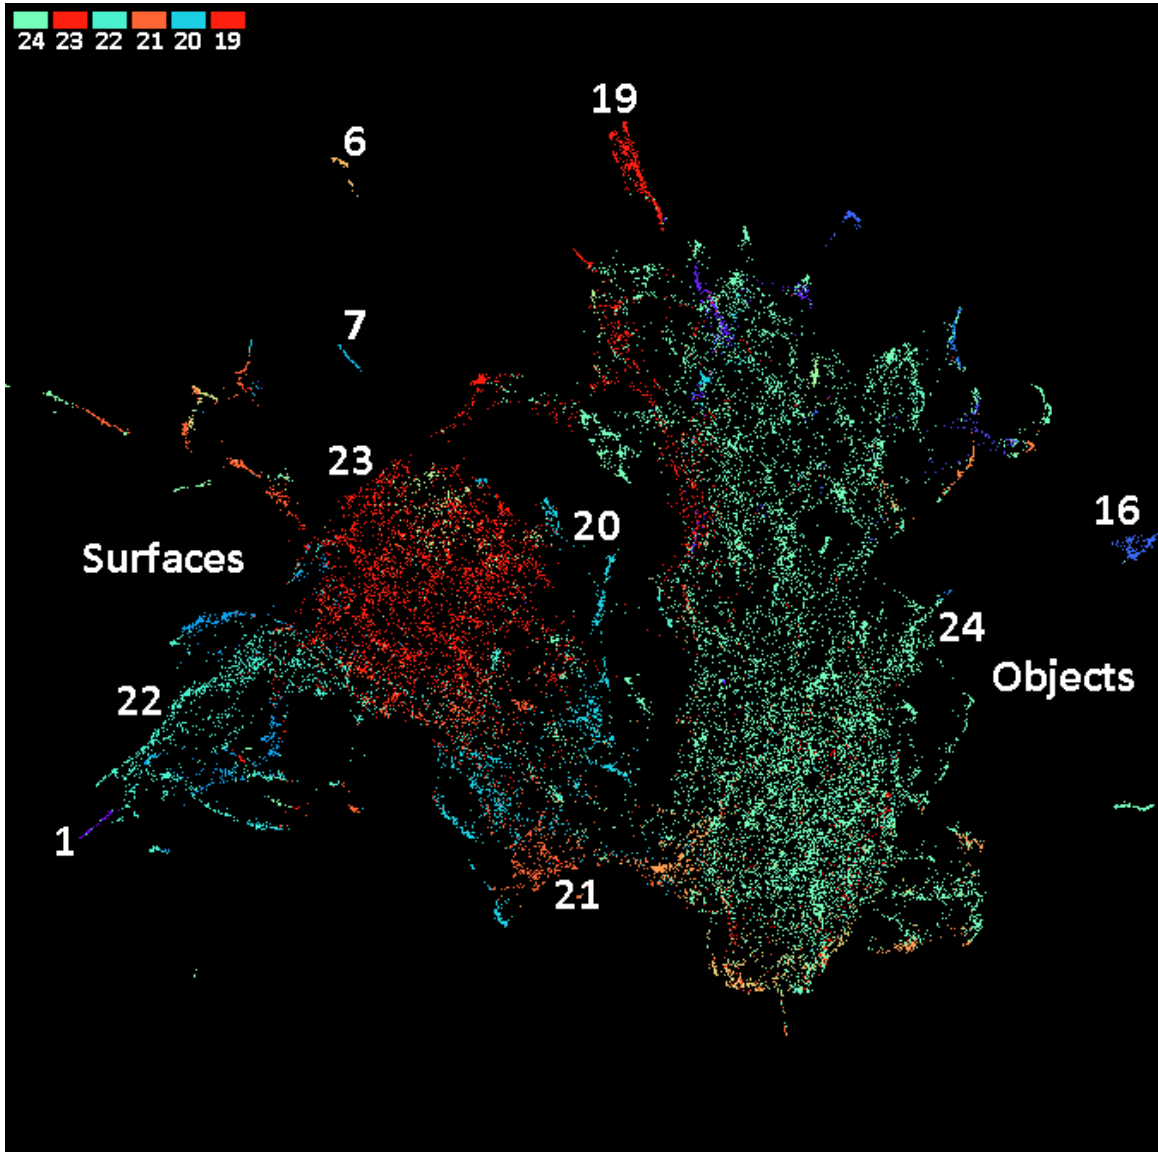


Figure 4.18: UMAP visualization of the ALOT ALOT-WCNN feature space. Minimum distance of 0.10 and 15 neighbors were used.

We observe a similar structure compared to the one we saw by the TCNN, except the split between the surfaces and object categories is more pronounced. On the right side, we have almost all the object classes captured by cluster 24 which appears to be an extension of cluster 28 from the TCNN experiment. On the left side, Clusters 20 to 23 make up the majority of the surface classes. Cluster 23 extends into a few object regions containing lattices. The other clusters are small and isolated within the projection. We observe cluster 1

from the TCNN experiment containing square grids is now grouped with cluster 24 despite the visible separation in the projection. After clustering with a noise threshold of 0, we found better separation in both the surfaces and objects with a similar cluster size distribution and UMAP visualization to the TCNN experiment. This separation was able to isolate the square grids cluster from the TCNN experiment indicating the ALOT-WCNN texture representation does not isolate unique representations from others as well. In general, the WCNN splits ALOT into fewer, larger clusters with less textural homogeneity. This might be accredited to a stronger focus on frequency analysis which works better at segmenting surfaces, but not necessarily as well on the shapes of objects.

Unlike Kylberg, we observed it was more difficult to break ALOT into finer clusters with a couple classes, however we did not see individual clusters forming due to intra-class variation. We suspect the primary reason ALOT is much harder to cluster is due to the number of classes. We note that we are using unsupervised learning, so we expect this to occur especially in similar textured classes. When two classes are similar, their feature representations can closely neighbor each other to the point that FDC observes their neighborhood as a single cluster. We saw this behavior in Kylberg with classes Stone1 and Stone2. If many classes neighbor each other, then by transitivity FDC may label a chain of clusters as a single elongated cluster. We noticed a glimpse of this behavior in the Kylberg-WCNN feature space from Figure 4.10. Having demonstrated the utility of our approach on Kylberg and ALOT using FDC and UMAP, we analyze the LPBF data in the same way in the next section.

4.2 LPBF In-situ Analysis

In the following sections, we cluster the recoat post-spread (4.2.1), recoat post-melt (4.2.2), thermal tomography (4.2.3), LWIR (4.2.4), and spatter (4.2.5) data. We repeat this clustering per experiment for comparison. Our goal is to determine which textural representations can isolate individual categories of process phenomena captured in each sensor modality. When clustering, we typically select a noise threshold of 0 to obtain the maximum number of clusters possible because we can always group similar clusters together manually. If a low noise threshold produces too many clusters to manage, we select a higher noise threshold. When results across experiments are similar, we limit discussions to highlight differences or reinforce earlier claims.

4.2.1 Recoat Post Spread

The recoat post-spread data consists of powder spreads from each coupon on each layer. It is known this modality reveals the presence of recoat streaking, recoat hopping, material protrusion, and soot. The textural representation of recoat streaking and hopping will likely be similar as they often show up as lines in the powder spread. Similarly, material protrusion and soot appear as small blobs. In this section, we perform clustering to group similar powder spreads. The benefit of clustering this data would be realized if powder spreads with the same roughness and anomalies group together cleanly. We begin by assessing how many clusters the recoat post-spread feature space can break into while varying the noise threshold of FDC. We display the number of clusters per threshold for each experiment in Figure 4.19.




Figure 4.19: Plot of the average number of clusters the post-spread data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.

The variation in the number of cluster between experiments indicates the feature set and model have an influence. When the noise threshold is very high, most experiments settle on the same number of clusters implying there may be a similar global clustering in each feature space. We dive into a clustering from each experiment.

Kylberg-TCNN

Starting with Kylberg-TCNN, we cluster the recoat post-spread feature space a noise threshold of 0 which produces 39 clusters. We relabel these clusters sorted by size and report them in Figure 4.20.






Figure 4.20: Recoat post-spread cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.


From this graph, we notice the post-spread data is breaking up into 8 significantly large clusters 32 to 39 and 31 small clusters. To find out why, we inspect each cluster by randomly sampling and describing the powder spreads from a textural perspective in Figure 4.21.




(1) Fine; bright corner/edge




(2) Smooth; large corners/edges




(3) Gradients; small dark corners; few large dark spots




(4) Smooth; large dark corners




(5) Gradients; small dark corners




(6) Fine; small bright corners; few dark spots




(7) Fine; small bright/dark corners/spots; lines




(8) Clumpy; shiny border




(9) Fine; large etched corners




(10) Coarse; few dark spots




(11) Smooth; large corners and spots




(12) Fine; tiny bright corners and dark spots




(13) Smooth; dark corners




(14) Smooth; Large bright/dark corners




(15) Smooth; large corner/edges




(16) Gradients; several small bright spots



(17) Fine; bright/dark corners; lines




(18) Smooth; bright corners/edges; small dark spots



(19) Smooth; bright/dark corners

(20) Fine; bright corners/edges; spots


(21) Coarse; tiny bright corners; several dark spots



(22) Fine; bright/dark corners; few lines and spots

(23) Smooth; shiny border


(24) Fine; small bright corners and dark spots



(25) Fine; bright spots; small dark spots

(26) Gradients; dark corners

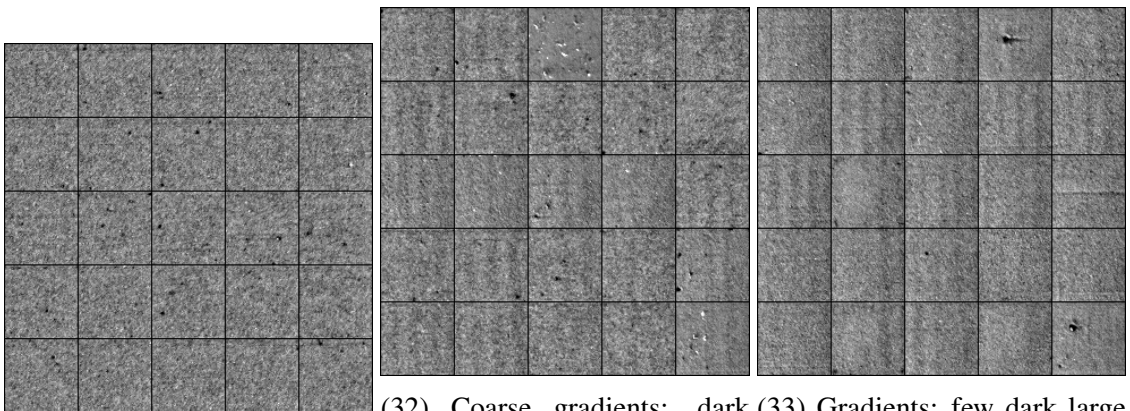
(27) Smooth; bright corners; few dark spots



(28) Gradients; several small spots

(29) Fine; large bright/dark corners; lines


(30) Fine; small bright/dark corners



(31) Coarse; dark spots

(32) Coarse gradients; dark spots


(33) Gradients; few dark large spots



(34) Fine gradients; lines; few dark spots

(35) Fine gradients; few spots

(36) Fine; lines; few dark spots,



(37) Coarse; several dark spots (38) Fine; few small spots (39) Fine; lines; dark spots

Figure 4.21: Normalized images sampled from each cluster in the recoat post-spread Kylberg-TCNN feature space. We describe the textural roughness of the powder spread followed by textural descriptions of other defects.

The Kylberg-TCNN texture representation segmented major and minor textural differences among the post-spread data. We observe most of the smaller clusters differ from the larger clusters by the presence of coupon corners protruding up through the powder layer. Among these smaller clusters, they segment apart by corner size and other phenomena in the powder.

This clustering demonstrates the roughness of powder spreads is a strong feature for segmentation because nearly all clusters are homogeneous in this manner, especially the large ones. Powder spreads segment depending on whether they are smooth, fine, coarse, or have gradients. These gradients are created by uneven lighting and reflections across the build plate. In some samples, the striped gradients hint at a subtle recoat hopping phenomena.

In general, most clusters are noisy in that they cannot isolate powder spreads sharing the exact same combination of features. We give two examples. Cluster 36 contains fine powder spreads, but the presence of lines and spots is not consistent. Similarly, cluster 14 contains smooth powder spreads with bright corner protrusion, but some include soot while others do not. We hypothesize the presence of some features (i.e. roughness of powder) may be a priority for clustering over others despite feature normalization.

We observe larger clusters 36 and 39 contain many samples with dead pixel residue lines.

This is sensor damage, not a process anomaly; however they do resemble the appearance of recoat streaking anomalies. These clusters suggest recoat streaking could be isolated if these lines were created from this phenomenon. In light of that, we do see pure recoat streaking in cluster 34. While these instances are more subtle, they are not grouped with dead-pixel residue lines.

Most clusters contain a mix of process phenomena with cluster 38 having the least. For a better interpretation of these clusters from the side of LPBF, we categorize them by the most common process phenomena from their powder spreads in Table 4.2. This labeling saves times, but sacrifices accuracy.

Phenomena	Clusters
Minor corner protrusion	1, 3, 5, 6, 7, 12, 17, 18, 19, 20, 21, 22, 24, 25, 28, 29, 30
Major corner protrusion	2, 4, 9, 11, 13, 14, 15, 27
Border protrusion	8, 23
Spot material protrusion	7, 9, 11, 16, 17, 28, 35, 39
Severe soot	37
Recoater Streaking	3, 7, 17, 29, 34, 36, 39

Table 4.2: Table of LPBF phenomena seen in Kylberg-TCNN recoat post-spread clusters.

Having analyzed the images from each cluster, we move onto the analysis of the spatial connection between clusters. To illustrate this, we color and label a UMAP projection of this feature space by cluster number in Figure 4.22.




Figure 4.22: UMAP visualization of the Recoat Post-Spread Kylberg-TCNN feature space colored by cluster number. Minimum distance of 0 and 50 neighbors were used.

We observe there are many sparse regions in the projection corresponding to most of the smaller clusters containing protruding coupon corners. The larger clusters 31 to 39 without corner protrusion compose the bulky regions. Near the center of the projection, there is a smaller cluster primarily containing powder spreads from many coupons with severe material protrusion, large soot, and recoat streaks. We show image samples from this cluster in Figure 4.23.




Figure 4.23: Normalized images sampled from severe defects cluster in the recoat post-spread Kylberg-TCNN feature space UMAP projection.

This cluster is visually isolated in the projection, but FDC grouped it with cluster 39. We believe the reason is textural overlap. The tail of cluster 39 transitions into spreads with small bright corners and material protrusion. The left side of the severe defects cluster starts with smaller bright spots and quickly transitions to the other defects mentioned earlier. In many samples, several features such as small bright/dark spots and lines are present simultaneously which contributes to the textural overlap.

We turn our attention to the bulky regions in the projection. Upon further inspection, we find these regions segment by coupon. To demonstrate this, we color the projection by coupon number in Figure 4.24.




Figure 4.24: UMAP visualization of the Recoat Post-Spread Kylberg-TCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.

The separation is clear and implies the physical location of the coupon on the build plate influences its post-spread image which in turn is reflected in its textural representation. We conclude the most likely cause is non-uniform lighting produced by various reflections and sources within the machine environment. Coupons physically located in the same row on the build plate tend to neighbor each other in the projection. This neighboring is tighter or entangled in the projection among coupons in the left and center columns. See coupon

pairs 1 and 2, 4 and 5, and 7 and 8. Coupons 4 and 5 are more distinct due to differences in their dead pixel residue.

While we suspect lighting plays a role, the airflow across the build plate may be too. Coupons 7, 8, and 9 reside in the front row and have much less soot in their spreads. When soot is carried away by the airflow, some of it may fall back onto the other coupons. After inspecting hundreds of samples from each coupon, we observe powder spreads from coupons in the first row have much less soot on average compared to the coupons in the proceeding rows. We also noticed the right column of coupons 3, 6, and 9 have less soot than their respective row-neighbors indicating the airflow may not be uniformly distributed. This concludes our analysis of the Kylberg-TCNN experiment. While this experiment produced the most number of clusters, we look at clusters from other experiments for comparison; however in the proceeding experiments, the analysis methods and most of the results are similar, therefore we limit our discussions to pointing out a few similarities and differences.

Kylberg-WCNN

Proceeding in the same way, we cluster the recoat post-spread Kylberg-WCNN feature space with a noise threshold of 0 which produces 22 clusters. We relabel clusters sorted by size and report them in Figure [4.25](#).





Figure 4.25: Recoat post-spread cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.


While we still have many small clusters, the 3 largest clusters show a significant jump from the other larger clusters. The sizes of those clusters are much higher than those previously implying this clustering may have less textural homogeneity. To verify this, we inspect randomly sampled images from each cluster and show them in Figure 4.26.



(1) Fine; bright corners

(2) Smooth; bright corners


(3) Smooth; large corners



(4) Smooth; corner/edges

(5) Fine; small bright corners;
few spots


(6) Clumpy; shiny border




(7) Fine; bright corners

(8) Fine; bright corners/edges;
few spots


(9) Fine; small bright corners;
spots




(10) Smooth; large etched corners




(11) Smooth; shiny border




(12) Fine; bright/dark corners; lines




(13) Fine; small dark corners; lines; large spots




(14) Coarse; bright/dark spots




(15) Smooth; bright corners



(16) Fine; small corners; large spots



(17) Fine gradient; lines



(18) Fine gradient; few spots

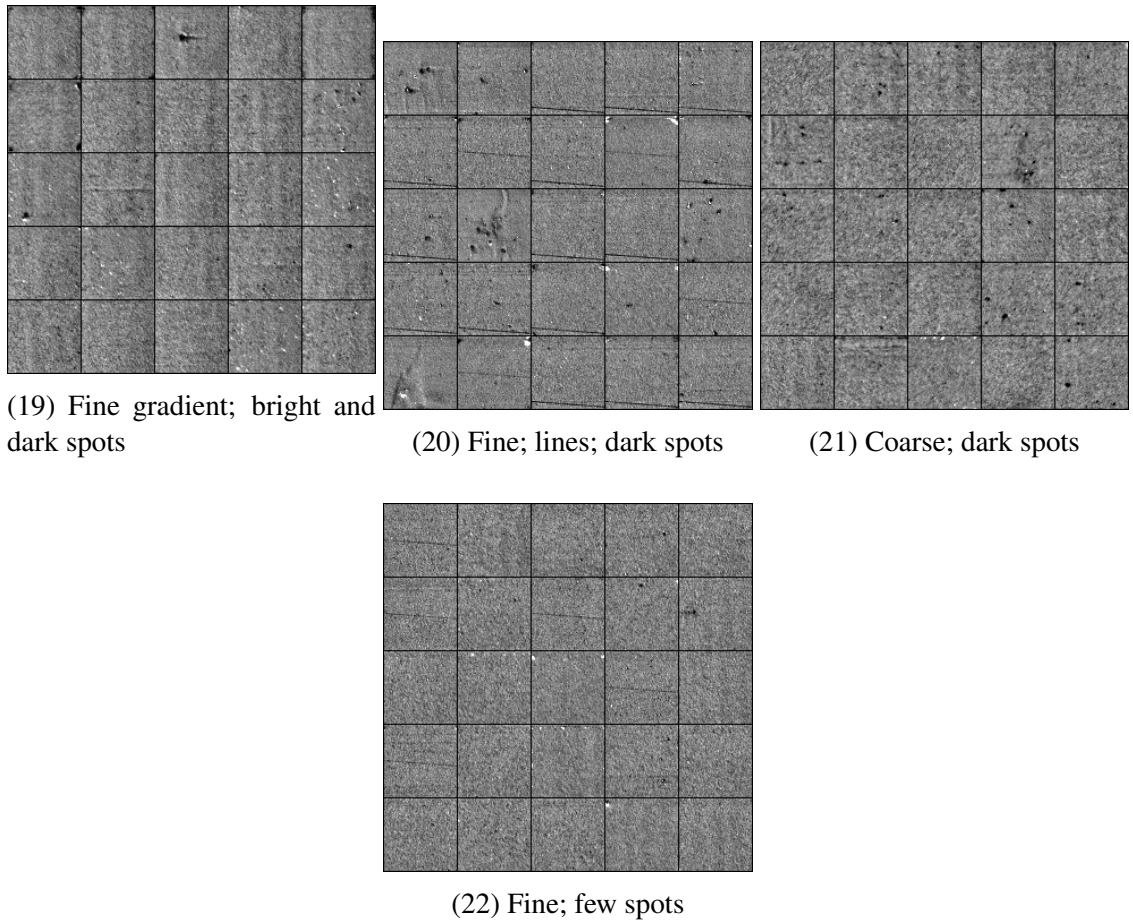


Figure 4.26: Normalized images sampled from each cluster in the recoat post-spread Kylberg-WCNN feature space.

Despite having much fewer clusters, many of them are texturally similar to those in Kylberg-TCNN implying the model architecture only provides local changes to the feature space. We give two examples of clusters shared across the Kylberg experiments. Cluster 17 contains a similar distribution of samples in Kylberg-TCNN cluster 34. They both describe fine powder spreads with subtle recoat streaking lines from coupon 6. Cluster 6 is identical to Kylberg-TCNN cluster 8. Both clusters are small containing powder spreads of the same significant coupon border protrusion.

Many small and some large clusters have merged together from Kylberg-TCNN. To see this better, we color UMAP projections of this feature space by cluster and coupon in Figures 4.27 and 4.28, respectively.

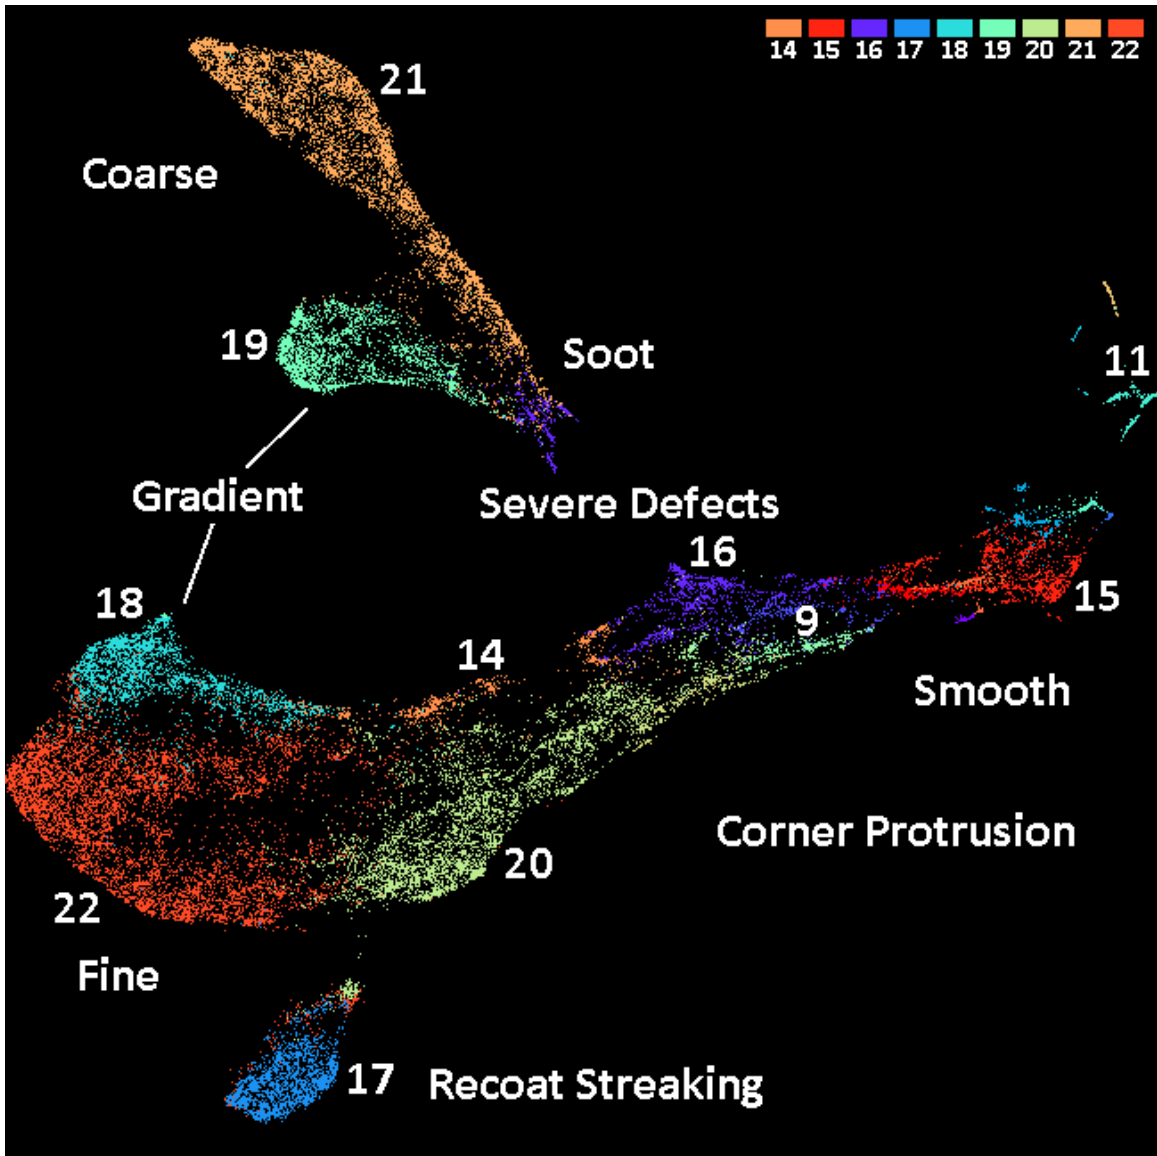


Figure 4.27: UMAP visualization of the Recoat Post-Spread Kylberg-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.




Figure 4.28: UMAP visualization of the Recoat Post-Spread Kylberg-WCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.

From these plots, we observe the geometry of this projection differs from the Kylberg-TCNN one, but the topology is very similar in that one end contains bulky regions that are coupon-homogeneous while the other end is reserved for corner protrusion. We observe coupons built in the right column of the build plate are still distinctly isolated in the projection. These observations are reflected in the clustering too, but we highlight some differences. Coupons 1 and 2 containing coarse powder spreads are better split in the pro-

jection and are grouped together into cluster 21. It is a similar case for coupons 4, 7, and 8 containing fine powder spreads in cluster 22. Many of the smaller Kylberg-TCNN clusters containing similar corner protrusion are now grouped into a broader cluster 15. We observe cluster 16 is at the same intersection between the two major clusters indicating it isolates the severe defect cluster found in the Kylberg-TCNN projection. This an improvement in clustering because the isolation of these defects supports anomaly detection. Clusters 19 and 21 are no longer split, which is unfortunate because some of their children clusters, such as Kylberg-TCNN cluster 16 containing speckled material protrusion, were more texturally homogeneous.

Overall, the WCNN produced a recoat post-spread feature space that is generally more connected which is seen by fewer clusters and a more contiguous projection. The textural homogeneity of most clusters was lower, with the redeeming exception of cluster 16 isolating most of the severe defects.

ALOT-TCNN

Next we consider experiments using features learned from ALOT. We cluster the ALOT-TCNN feature space with a noise threshold of 0 which yields 25 clusters. We relabel clusters sorted by size and report them in [Figure 4.29](#).





Figure 4.29: Recoat post-spread cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.


The distribution of cluster sizes is more uniform for clusters 20 to 23 and larger clusters 24 and 25. The remaining clusters gradually decay in size. This distribution differs more from the Kylberg experiments in that the larger clusters are smaller and vice versa implying large clusters are splitting and smaller clusters are merging. Both effects are desirable for isolating broad, texturally homogeneous clusters. We randomly sample images from each cluster and label them in Figure 4.26.



(1) Fine; bright edges

(2) Clumpy; shiny border


(3) Fine; bright corners



(4) Smooth; shiny border

(5) Smooth; large corner edges


(6) Coarse gradients; dark spots




(7) Fine; small corners

(8) Fine gradients; small corners


(9) Smooth or fine gradients; dark corners and spots




(10) Coarse; few dark spots




(11) Coarse; large dark spots




(12) Fine; corners, lines, and raised spots




(13) Fine; corners, lines, and raised spots




(14) Coarse; many large dark spots




(15) Smooth; large corners



(16) Fine; large corners



(17) Coarse; few bright and dark spots



(18) Fine; small bright corners and spots

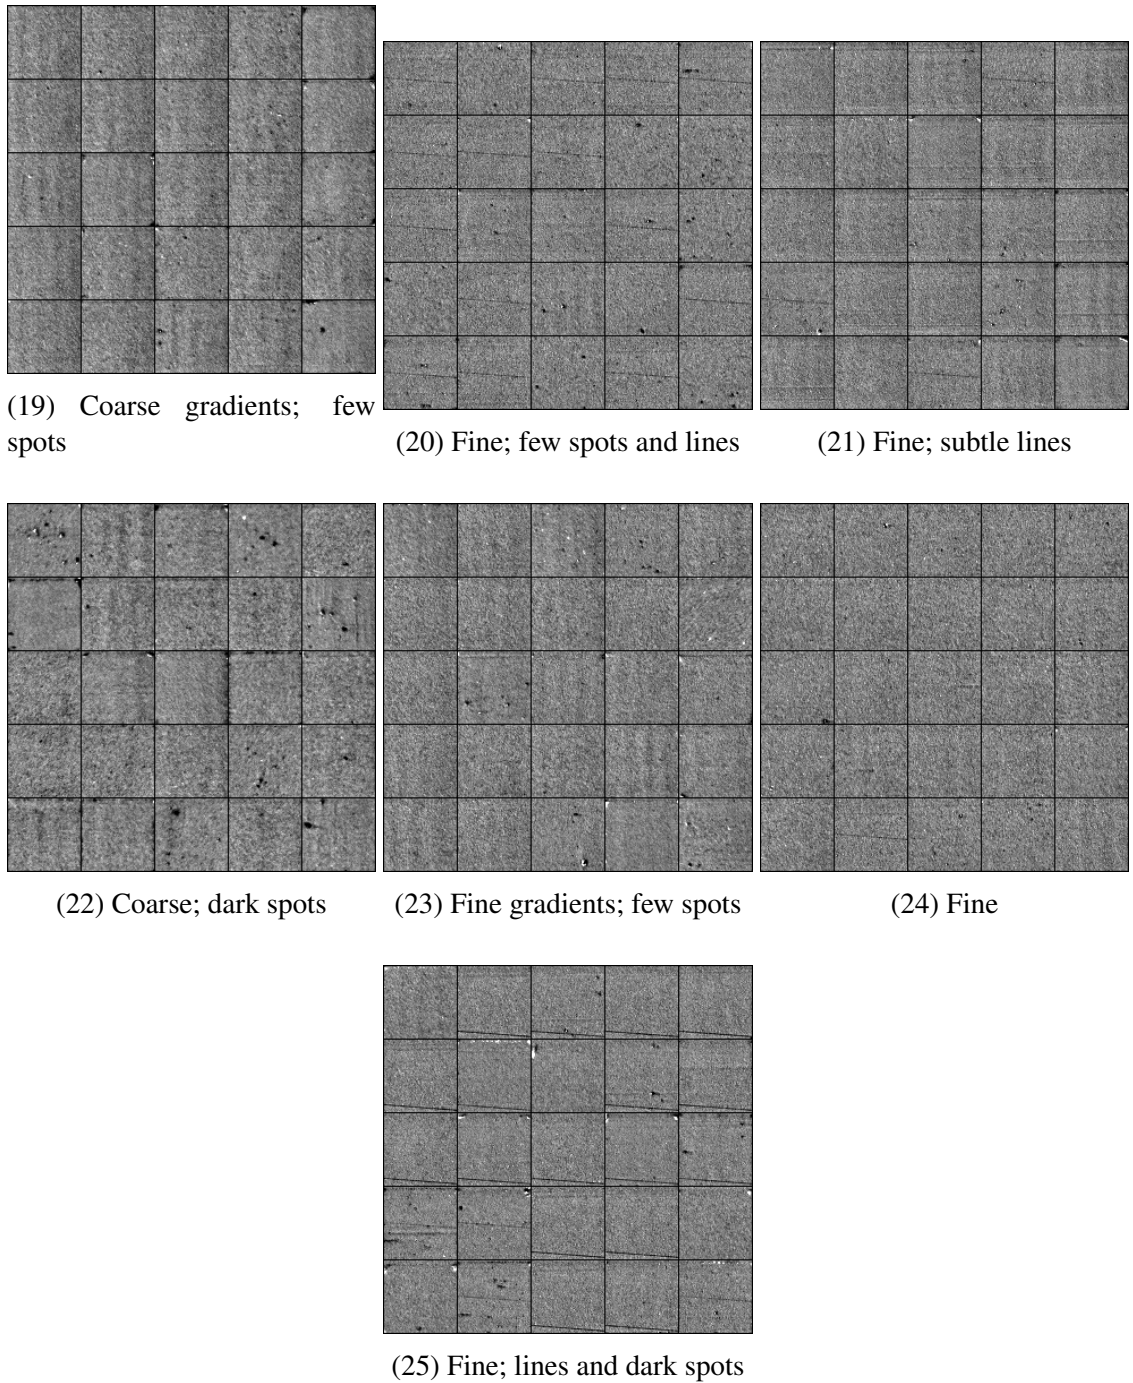


Figure 4.30: Normalized images sampled from each cluster in the recoat post-spread ALOT-TCNN feature space.

Once again, coupon border material protrusion isolates well with clusters 16 and 18 grouping larger and smaller corners together. Smaller clusters 10, 11, 14, and 17 containing coarse powder spreads differ in the amount of soot implying the cluster containing coupons

1 and 2 has split. The severe defect cluster is not well isolated like in the Kylberg-WCNN experiment. We see similar large clusters 20 to 25 likely corresponding to coupons as seen in previous experiments. To illustrate these clusters further, we color and label the UMAP projection of this feature space by cluster and coupon in Figures 4.31 and 4.32, respectively.

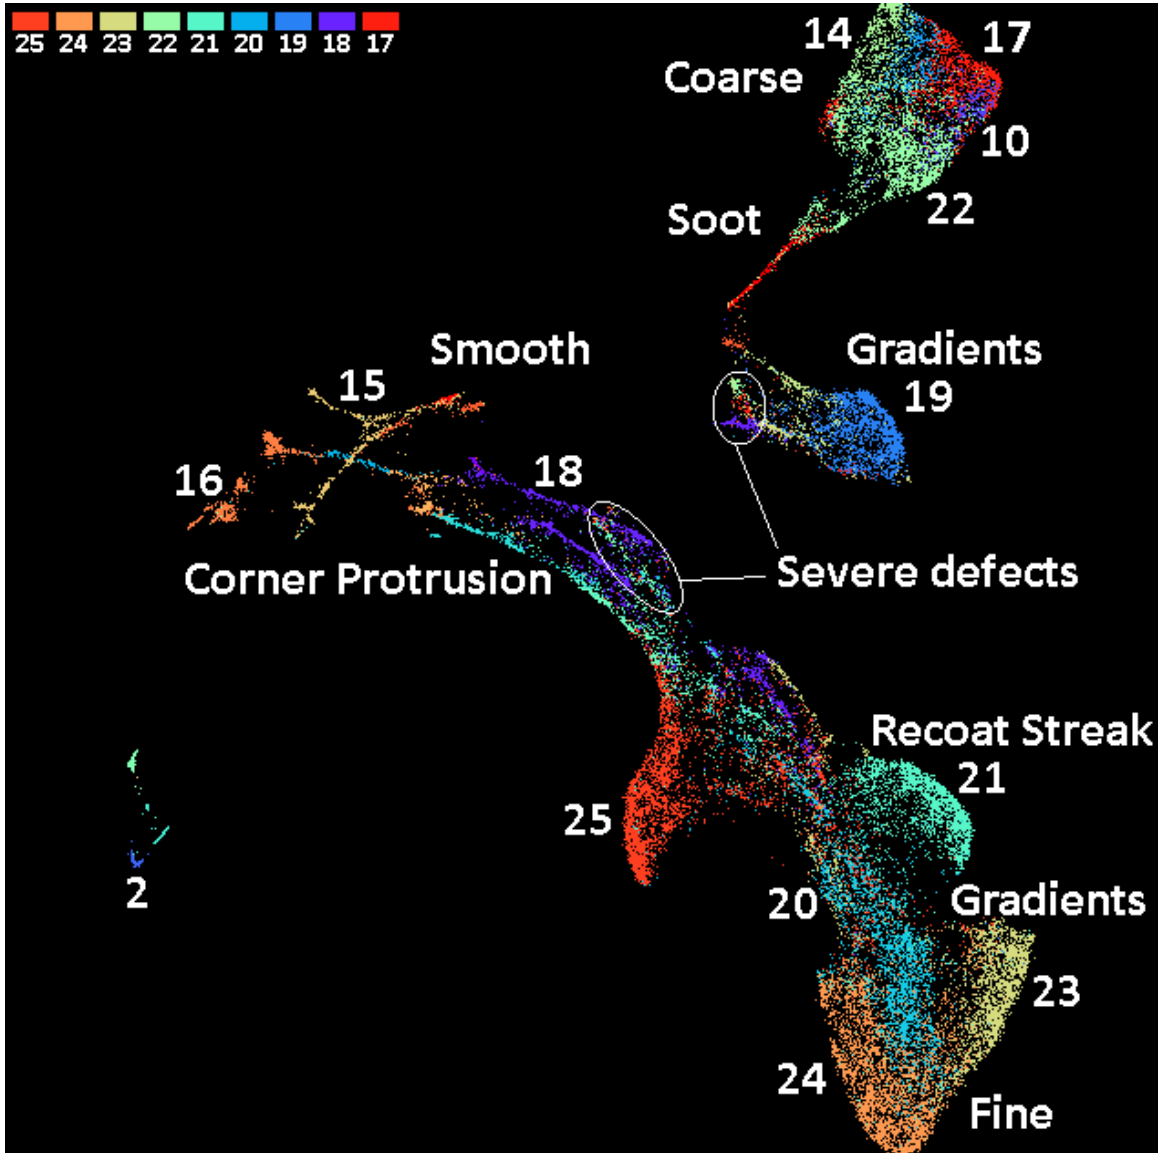


Figure 4.31: UMAP visualization of the Recoat Post-Spread ALOT-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.

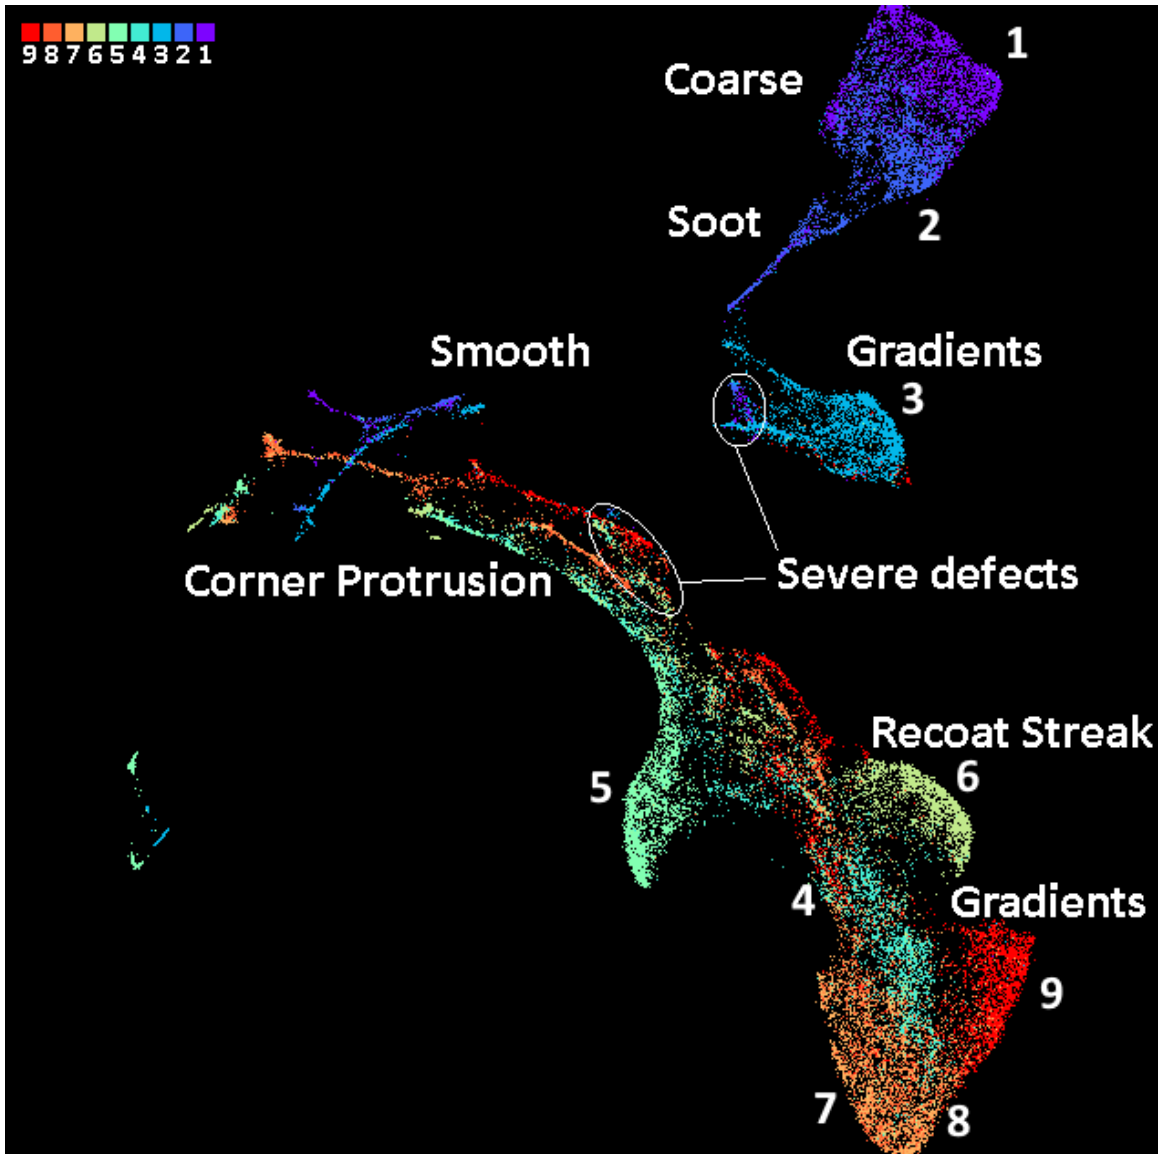


Figure 4.32: UMAP visualization of the Recoat Post-Spread ALOT-TCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.

The projection shape is different at a local scale, but overall the experiments tend to settle on the same topology that separates powder spreads by the presence of a corner protrusion. The separation of coupons is still clear, with the exception of the entangled coupons 7 and 8. From these colorings, the previously mentioned smaller clusters 10, 11, 14, and 17 are primarily splitting coupon 1. One difference that could be the cause is coupon 2 has more samples with gradients while coupon 1 does not. We notice the severe defect cluster has

split in the projection and clustering; however similar anomalies are isolated locally in the projection but are grouped into larger clusters with different phenomena. It is not clear why this is the case. Clusters 15, 16, and 18 isolate the different levels of coupon corner protrusion which helps generalize the clustering without sacrificing semantic homogeneity. A similar generalization was recognized in the Kylberg-WCNN experiment, but not in the Kylberg-TCNN experiment because it break ups into many more smaller clusters. It is becoming apparent that each of these experiments have their own strengths and weaknesses, but generally agree on many clusters.

ALOT-WCNN

Finally, we cluster the recoat post-spread ALOT-WCNN feature space with a noise threshold of 0 yielding 21 clusters. We relabel these clusters sorted by size and report them in Figure 4.33.





Figure 4.33: Recoat post-spread cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.

We see a similar distribution to the one from the ALOT-TCNN experiment, except instead of two equally large clusters, we have one large cluster with over twice the size. We


suspect the two large clusters from another experiment combined to produce cluster 21. To verify, we randomly sample images from each cluster and label them in Figure 4.34.



(1) Fine; corners; lines; spots

(2) Fine; corners; lines; spots


(3) Smooth; large corners



(4) Fine; bright corners

(5) Fine; lines; small corners; spots


(6) Fine; many lines; spots




(7) Coarse; dark and bright spots

(8) Smooth; shiny border


(9) Fine; small spots




(10) Smooth; bright/dark corners; spots




(11) Coarse gradients; spots




(12) Variety of border and corners




(13) Fine; bright corners; few spots and lines




(14) Fine; bright/dark spots




(15) Coarse gradients; few spots



(16) Fine gradients; lines; spots



(17) Fine gradients; lines



(18) Fine; lines; dark spots

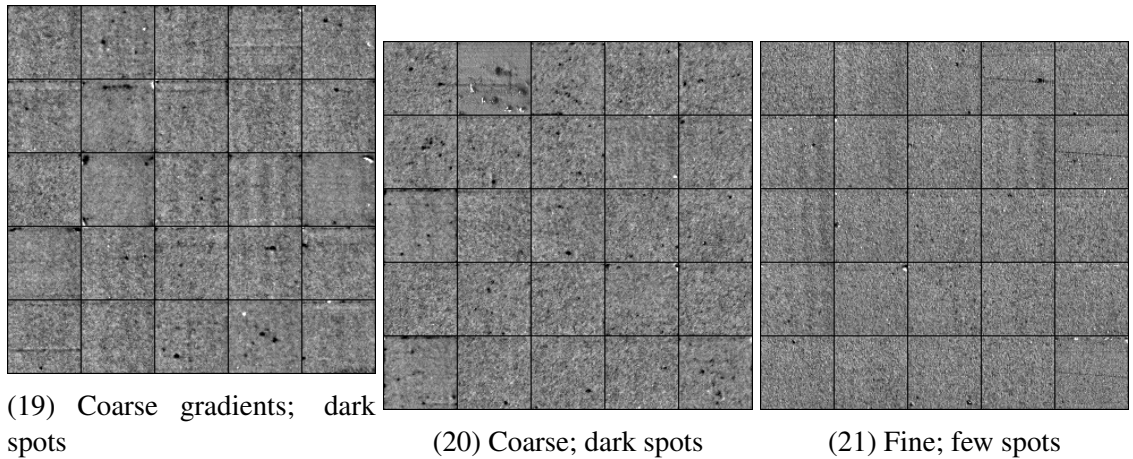


Figure 4.34: Normalized images sampled from each cluster in the recoat post-spread ALOT-WCNN feature space.

We see similar clusters from previous experiments; however we get different separation and merging. Regarding separation, cluster 7 was able to isolate abundant small bright spot material protrusion similar to Kylberg-WCNN cluster 16. This phenomenon has only been isolated in WCNN experiments implying frequency analysis is better at embedding this feature. Clusters 1, 2, 5, and 6 contain samples with more severe defects, but do not completely make up the severe defect cluster. Unlike the Kylberg-WCNN, the distinction between coupons 1 and 2 has been preserved. Regarding merging, corner protrusion clusters 8 and 12 are a bit noisy varying in levels of severity. Cluster 21 being much larger with a few subtle gradient samples implies the distinction between coupons 7, 8, and 9 has been lost. In the ALOT-TCNN experiment, the coupon 1 cluster containing coarse powder spreads with no gradients split into varying levels of soot, but this splitting is not realized by cluster 20. To emphasize these observations, we color the UMAP projection of this feature space by cluster and coupon in Figures 4.35 and 4.36, respectively.

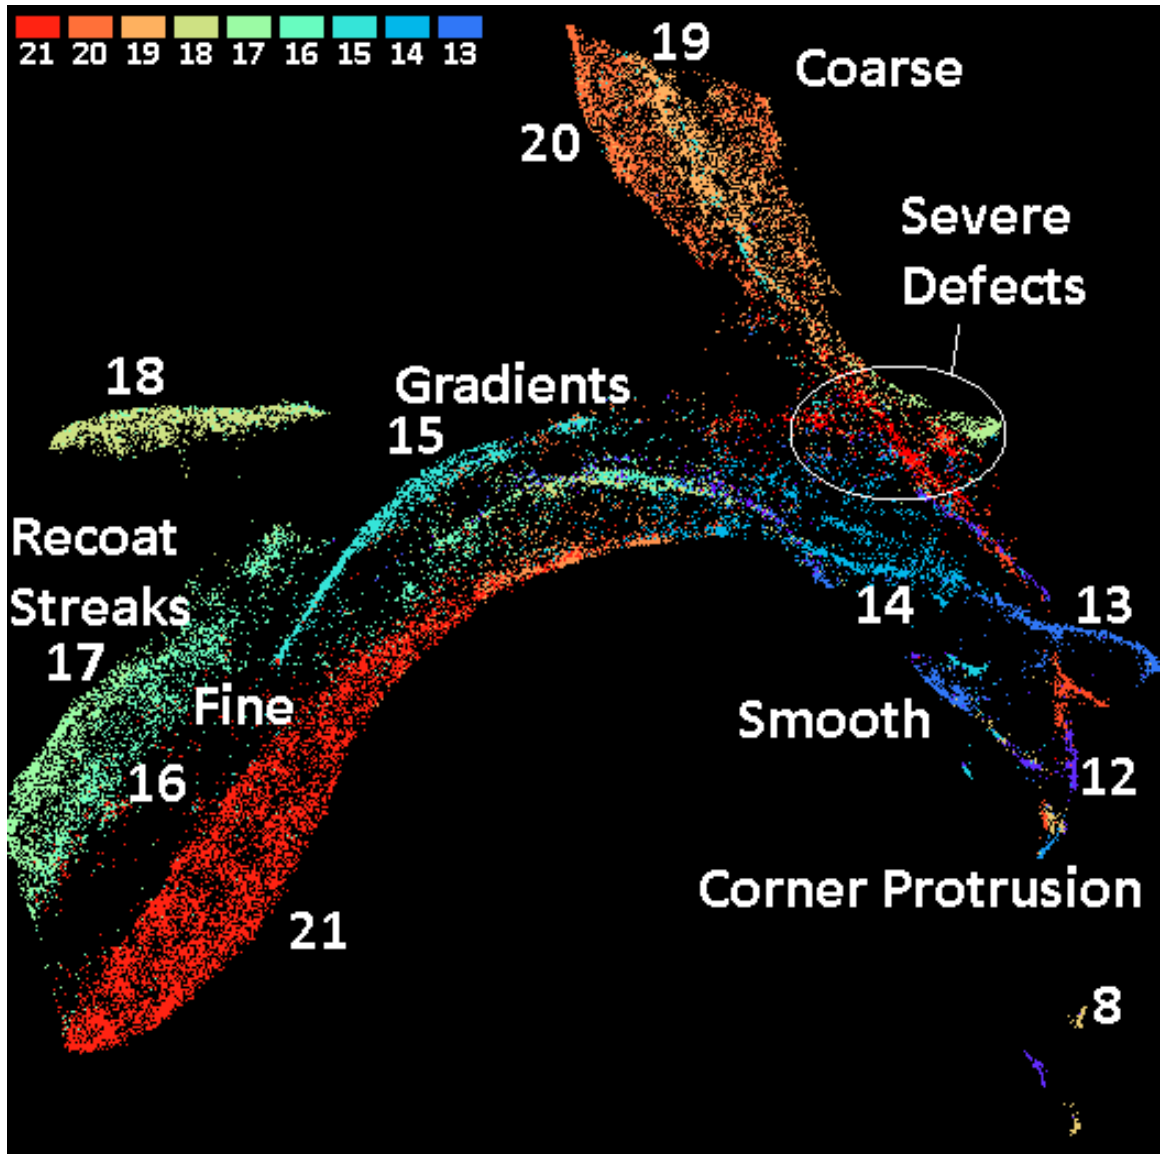


Figure 4.35: UMAP visualization of the Recoat Post-Spread ALOT-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.

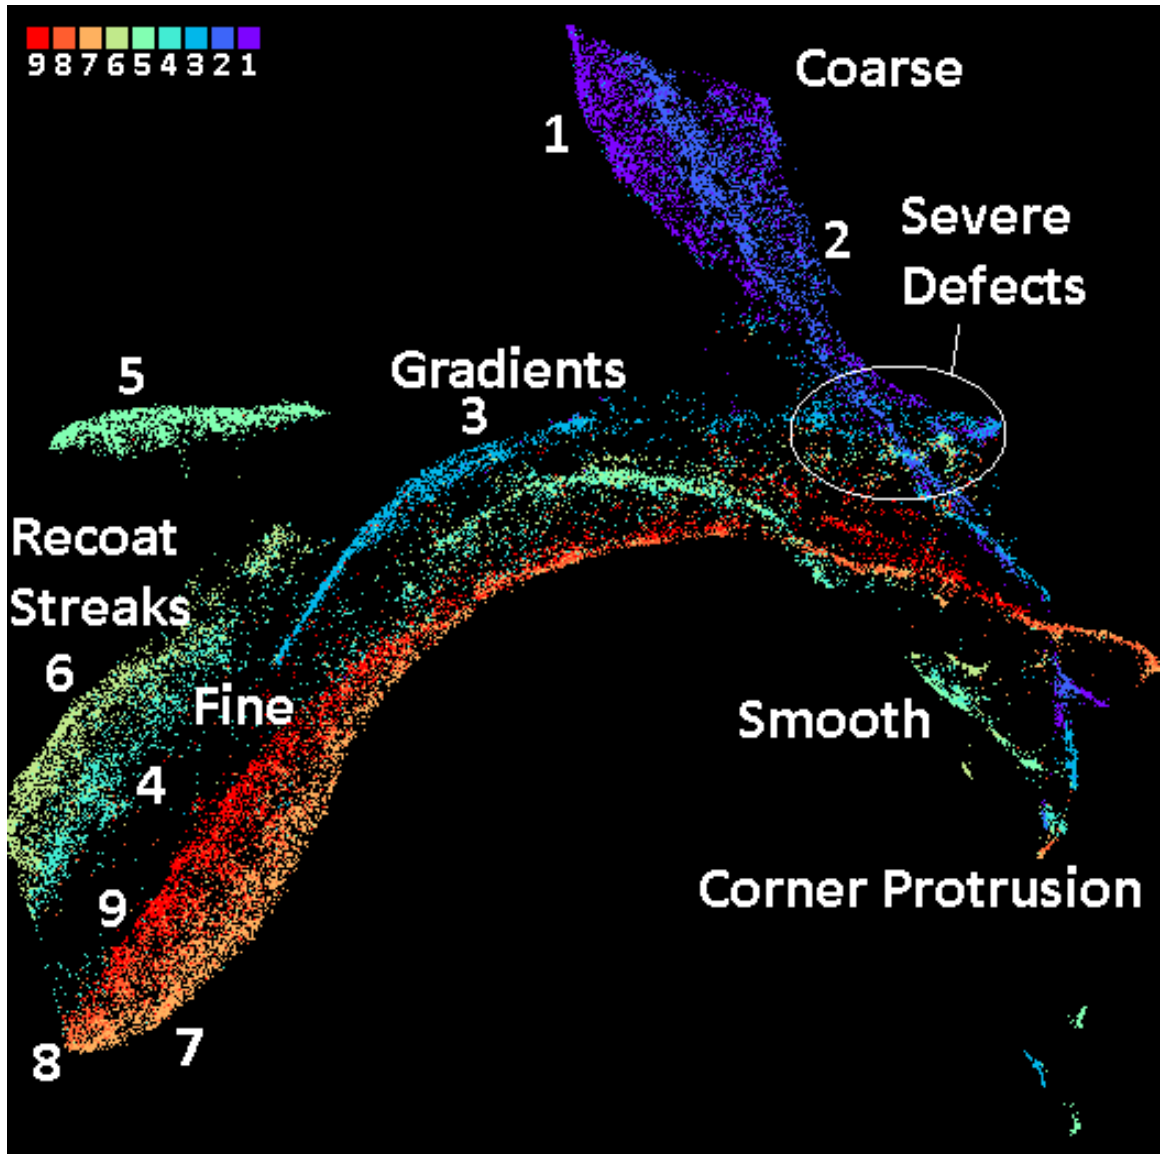


Figure 4.36: UMAP visualization of the Recoat Post-Spread ALOT-WCNN feature space colored by coupon number. Minimum distance of 0 and 50 neighbors were used.

This projection differs most from the other experiments. The coupon clusters are stretched alongside each other. Coupon 3 is a prime example of this. From the Kylberg experiments, coupon 6 breaks away leaving coupons 4 and 5 being close neighbors, but cluster 5 breaks away instead. A similar but more subtle observation is realized in the ALOT-TCNN projection. We suspect the difference in ALOT and Kylberg feature embedding must be the cause.

Coupons 7, 8, and 9 are tightly strung alongside another in the projection. It follows that they are all grouped into cluster 21. Coupons 1 and 2 group into clusters 20 and 19 but do not split further like in the TCNN experiments. Once again, at the tail of coupons 1 and 2, we find the severe defect cluster. Some categories of the defects are distinguished, such as speckled material protrusion in cluster 7, but most others end up being grouped within larger clusters. For example, cluster 14 primarily contains powder spreads with tiny bright corners, but it also captures spreads with several small bright material protrusion spots. The textural representation is likely similar and less sensitive to positioning.

Summary

In the recoat post-spread data, we effectively clustered roughness categories of powder spreads ranging from smooth, fine, and coarse. Among these, there were subcategories with material protrusion, recoat streaking, soot, and severe defects. UMAP projections isolated a region with severe defects; however this region rarely isolated well in the clustering. The recoat post-spread feature space UMAP projections were globally similar due to the high sensitivity to coupon positioning.

The influence of coupon positioning led to most powder spread images grouping together if they came from the same coupon. There were several driving forces behind this bias. Non-uniform lighting across the columns revealed subtle gradients only on right column coupon powder spreads. The focus of the camera varied along the rows thereby producing coarser powder spread imagery for coupons in the back row. Sensor damage varied for coupons 4 and 5 thereby making them more distinct. Additionally, soot was more common in the back row powder spreads due to higher turbulence in the argon cross-bed airflow. These biases were confirmed separately by FDC and UMAP across all experiments. This unfortunately makes most defect labeling inaccurate and spatially dependent.

We consider a few modifications for the future. Corner protrusion was a significant feature and sometimes confused with spot material protrusion. The distinction relies on knowing

the part geometry; but if images were cropped to exclude coupon borders before clustering, the internal areas might cluster better. This could be the case in general if the window size is reduced as there would be less textural overlap and anomalies within the part geometry would have a larger footprint in the field of view. Mitigating the effect of lighting biasing the texture distribution across the build plane is a high priority. We propose two techniques. A preprocessing step such as localized histogram equalization or lighting calibration could reduce global lighting gradients but may not be sufficient to remove local lighting artifacts. Positioning light sources in a more uniform arrangement may provide more equal illumination across the build plane.

Among these experiments, we found clustering prioritized the roughness of the powder and whether coupon corner protrusion was present. The TCNN experiments produced sparser projections and more clusters that better separated samples by levels of soot and corner protrusion. The WCNN experiments produced highly-connected projections and fewer clusters with less textural homogeneity. Each experiment had pros and cons; however we reserve the Kylberg experiments for further research as the Kylberg-TCNN had the most separation while the Kylberg-WCNN isolated the severe defects cluster.

In summary, the recoat post-spread data was heavily biased and segmented by roughness category and the presence of material protrusion. Our goal of isolating soot and severe defects was barely achieved with the few clusters capturing these phenomenon exhibiting many false positives. Since the recoat post-melt data is captured by the same sensor, we suspect the lighting will have a similar influence.

4.2.2 Post Melt

The recoat post-melt data consists of melted powder from each coupon on each layer. It is known this modality can reveal characteristics of the melted material surface roughness. Surfaces can appear as glossy, shiny, etched, grooved, coarse, bumpy, fine, and/or rough. Spatter can land on the powder spread introducing sporadic bumps on the melted surface making it rough. These bumps can show up as material protrusion in the proceeding powder spread. When a melted surface is etched, it is revealing the laser scan lines that melted it. Glossy surfaces can make the part shine but also experience texture washout where the entire surface reflects too much light causing over-saturation. This leads to a trivial texture where all the pixels are the max value, so we exclude these instances from our analysis. In this section, we perform clustering to group similar melted surfaces. This clustering would be beneficial if broad categories of surfaces can be isolated. We begin by assessing how many clusters the post-melt feature space can break into while varying the noise threshold of FDC. We display the number of clusters per threshold for each experiment in Figure 4.37.




Figure 4.37: Plot of the average number of clusters the post-melt data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.

The number of clusters varies a lot depending on the experiment. In the post-spread

recoat analysis, we found the Kylberg-TCNN representation produced the most clusters by a large margin; however the opposite holds for this data demonstrating one experiment may not be sufficient for clustering all datasets. Additionally, we cannot generally conclude a single model nor dataset will produce a sparser feature space resulting in more clusters, thereby re-emphasizing the need for multiple experiments to obtain a good clustering. We note the ALOT-TCNN experiment produces the most clusters at a noise threshold of 0, but it is also the most sensitive to the random seed indicating there is more fluidity among clusters. At higher noise thresholds, the Kylberg-WCNN and ALOT-TCNN produce nearly the same number of clusters implying their clusters may be similar. We dive into a clustering for each experiment.

Kylberg-TCNN

Starting with Kylberg-TCNN, we cluster using a noise threshold of 0 which produces 17 clusters, the fewest out of all experiments. We relabel these clusters sorted by size and report them in Figure 4.38.






Figure 4.38: Recoat post-melt cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.

At first glance, most of the data is segmented into two general categories, which can


be uninformative if clusters 16 and 17 are noisy or too broad; however we remark it is unlikely coupon positioning is affecting this clustering as we would expect to see more equally large clusters corresponding to each coupon. We look to smaller clusters 1 to 15 to provide us with more meaningful and cleaner separation. We inspect these clusters by randomly sampling and labeling their images in [Figure 4.39](#).




(1) Glossy; few bright specks; large bright corners/edges




(2) Solid; few very dark spots




(3) Glossy etching; bright corners




(4) Coarse; uniformly bumpy




(5) Shiny; large gradients




(6) Smooth etching; raised corners/edges




(7) Fine etching; specks




(8) Smooth etching; raised corners/edges




(9) Smooth and glossy; few bright speck/corners/edges



(10) Glossy; large bright corner/edges (11) Smooth etching; bright spots (12) Glossy; bright specks/corners



(13) Smooth; several subtle specks; bright edges (14) Coarse and wavy (15) Smooth; rough specks; bright edges



(16) Rough and coarse (17) Rough

Figure 4.39: Normalized images sampled from each cluster in the recoat post-melt Kylberg-TCNN feature space.

In the larger clusters 16 and 17, we have an overlap of rougher surfaces making up the majority of the post-melt data. Cluster 16 differs from cluster 17 in that samples are finer

with smaller bumps and waves. Subcategories of these samples from these two clusters have yet to be observed in other experiments. Unlike in the post-spread analysis, smaller clusters do not segment from larger ones primarily by coupon corners but rather by the surface roughness. As expected, the smaller clusters are texturally homogeneous. We provide several examples of these. Clusters 2 and 5 contain shiny surfaces nearly washed out from over-saturation. Clusters 1, 10, and 12 capture glossy surfaces varying in the size of shiny corners and edges. Clusters 4 and 14 isolate coarse bumpy surfaces. Clusters 6 and 8 group smooth surfaces whose reflections indicate significantly raised edges. Clusters 11, 13, and 15 capture smooth surfaces varying in the size of shiny edges, corners, and severity of specks. Cluster 7 alone captures fine surfaces. To understand how these clusters are organized, we color a UMAP projection of this feature space by cluster number in Figure [4.40](#).

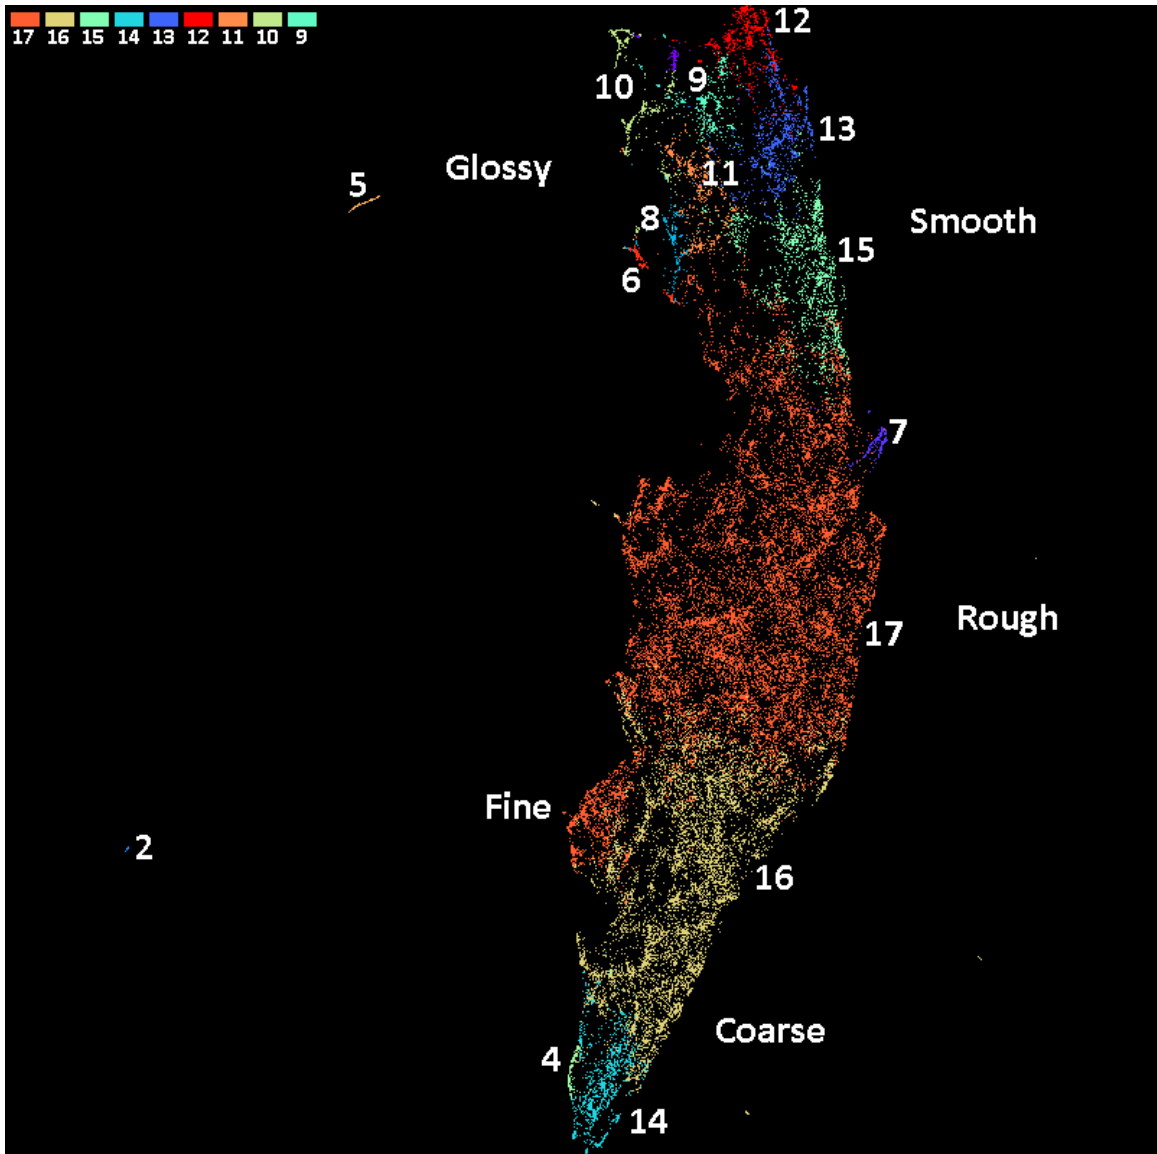



Figure 4.40: UMAP visualization of the Recoat Post-Melt Kylberg-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.


From this projection, we observe the feature space forms a single cluster containing nearly all the data. After running UMAP with fewer neighbors, we verify this is still the case likely due to high overlap in the feature space. Looking at the arrangement of these clusters, we find rough, coarse, and fine surfaces in clusters 17 and 16 towards the center. The bottom end transitions into coarse bumpy surfaces in clusters 4 and 14. On the left side of the rough surfaces region, we find several fine surfaces, but most of these grouped

into clusters 17 followed by 16. Interestingly, cluster 7 unexpectedly does not appear on the same side where the fine surfaces are located. Towards the upper end of the rough region we find smoother surfaces within clusters 11, 13, and 15 which eventually transition into glossy surfaces in clusters 9, 10, and 12. We observe clusters 6 and 8 contain shiny and smooth surfaces with rough corners and edges curling up which places them alongside the transition from rough to smooth surfaces. As clusters 2 and 5 are much different from other samples in that most of the surface is over-saturated, it follows that their samples are isolated far from others in the projection. Before moving onto the next experiments, we take a moment to assess the influence of coupon positioning and sensor damage by coloring the projection by coupon row, column, and index in [Figure 4.41](#).



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.41: UMAP visualization of the Recoat Post-Melt Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

The influence of coupon column positioning appears negligible as coupons from each column are uniformly distributed along the projection. This is not the case for the coupon row as samples spatially cluster together in the projection with the same row index. This is also illustrated by veins in the coupon index coloring. For instance, observe how back row coupons 1, 2, and 3 run along the right side of the projection together. This happens

similarly for the other coupons. The middle row coupon surfaces are influenced the most. These coupons compose most of the fine region while avoiding the coarse region in the projection. This impacts the clustering in that cluster 16 captures much fewer middle row surfaces; however rough surfaces in cluster 17 remain unaffected. The bias could be related to more light being concentrated over the middle row thereby casting minimal shadows and producing a grainy image. The recoat post-spread analysis revealed a similar phenomenon leading to this same reasoning.

Regarding dead pixel residue lines induced by prior sensor damage, they are barely visible in the samples and show little to no influence on the clustering and projection. We conclude the reason is related to the distribution of pixel values before normalization. Post-melt imagery typically has a wider range and higher standard deviation of pixel values that surpass the deviation of dead pixel residue from the mean. This prevents dead pixel residue from being mapped to the darkest values during normalization thereby making them stand out less and avoid being embedded as a significant feature in the textural representation. This concludes our analysis of this baseline Kylberg-TCNN experiment. We continue by pointing out similarities and differences in other experiments.

Kylberg-WCNN

Proceeding with the Kylberg-WCNN, We cluster the recoat post-melt feature space with a noise threshold of 0 to produce 36 clusters, which is over twice that in the baseline Kylberg-TCNN clustering. We relabel these clusters sorted by size and report them in Figure 4.42.






Figure 4.42: Recoat post-melt cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.


This distribution is very different in comparison to the Kylberg-TCNN cluster size distribution. It is more uniform in that the bigger clusters have split into several medium-sized ones. We can expect more textural homogeneity from these and the smaller clusters. We visually assess this intuition by inspecting randomly sampled images from each cluster in Figure 4.43.




(1) Grooved; Bright top edge




(2) Shiny; large gradients




(3) Shiny; large gradients




(4) Solid; few dark spots




(5) Coarsely fine




(6) Glossy; bright corners/edges




(7) Coarse and wavy




(8) Glossy etching; bright corners




(9) Glossy etching; bright edges




(10) Grooved and rough; dark edges/spots; bright




(11) Glossy, large shiny corners




(12) Smooth etching; raise corners/edges




(13) Glossy; large corner gradients




(14) Grooved and rough; dark edges/spots; bright




(15) Smooth etching; bright corners/edges




(16) Etched; small bright spots/edges/corners




(17) Smooth; subtle spots; bright corners




(18) Glossy; subtle spots; bright corners/edges




(19) Smooth; small specks; bright corners/edges




(20) Smooth; rough spots; small bright corners/edges




(21) Coarse and wavy




(22) Coarse; large specks




(23) Fine etching; specks; minor edges




(24) Coarse




(25) Fine etching; rough specks



(26) Smooth etching; rough specks




(27) Coarse



(28) Semi-rough; specks

(29) Finely coarse; specks


(30) Rough; large specks



(31) Rough, semi-coarse

(32) Smooth and fine; subtle specks; bright edges/corners

(33) Smooth; rough specks; bright edges



(34) Coarse, semi-rough

(35) Semi-rough; grooves and spots

(36) Semi-rough; bright specks

Figure 4.43: Normalized images sampled from each cluster in the recoat post-melt Kylberg-WCNN feature space.

From these cluster images, we observe a higher textural homogeneity not only for smaller clusters but for many larger ones too. Among the new additions of clusters, we ob-

tain more subcategories of surfaces. We give two examples. Clusters 1, 10, and 14 isolate bright rough surfaces that reveal grooves. Clusters 28 to 31 and 34 to 36 split rougher surfaces varying on levels of coarseness and surface features. We point out a couple instances of this clustering that are too specific. Cluster 1 captures near-identical grooved surfaces with a bright top edge when they could easily fit into cluster 14. Clusters 2 and 3 could be grouped as they capture the same shiny surfaces. Some clusters are capture near-identical groups but differ in a small feature. Cluster 6, 8, and 9 contain glossy samples, but vary in the severity of etching. Cluster 6 has none, while the etching in cluster 8 is sharper than in cluster 9.

In general, this finer clustering is superior to the one we saw in the Kylberg-TCNN experiment. We find it possible to provide a reasonably accurate labeling of clusters corresponding to general surface categories in Table 4.3. Note some clusters must be represented in multiple surface categories because clusters are not truly homogeneous.

Surface Category	Clusters
Glossy	6, 8, 9, 11, 13, 18
Smooth	12, 15, 17, 18, 19, 20, 26, 32
Fine	5, 16, 23, 25
Rough	28, 30, 34, 35, 36
Coarse	5, 7, 21, 22, 24, 27, 29, 31, 34
Etching	8, 9, 12, 15, 16, 23, 25, 26
Grooves	1, 10, 14
Saturated	2, 3, 4, 10, 14

Table 4.3: Table of surface categories seen in Kylberg-WCNN recoat post-melt clusters.

We compare projections next. We color the UMAP projection of this feature space by cluster in Figure 4.44.

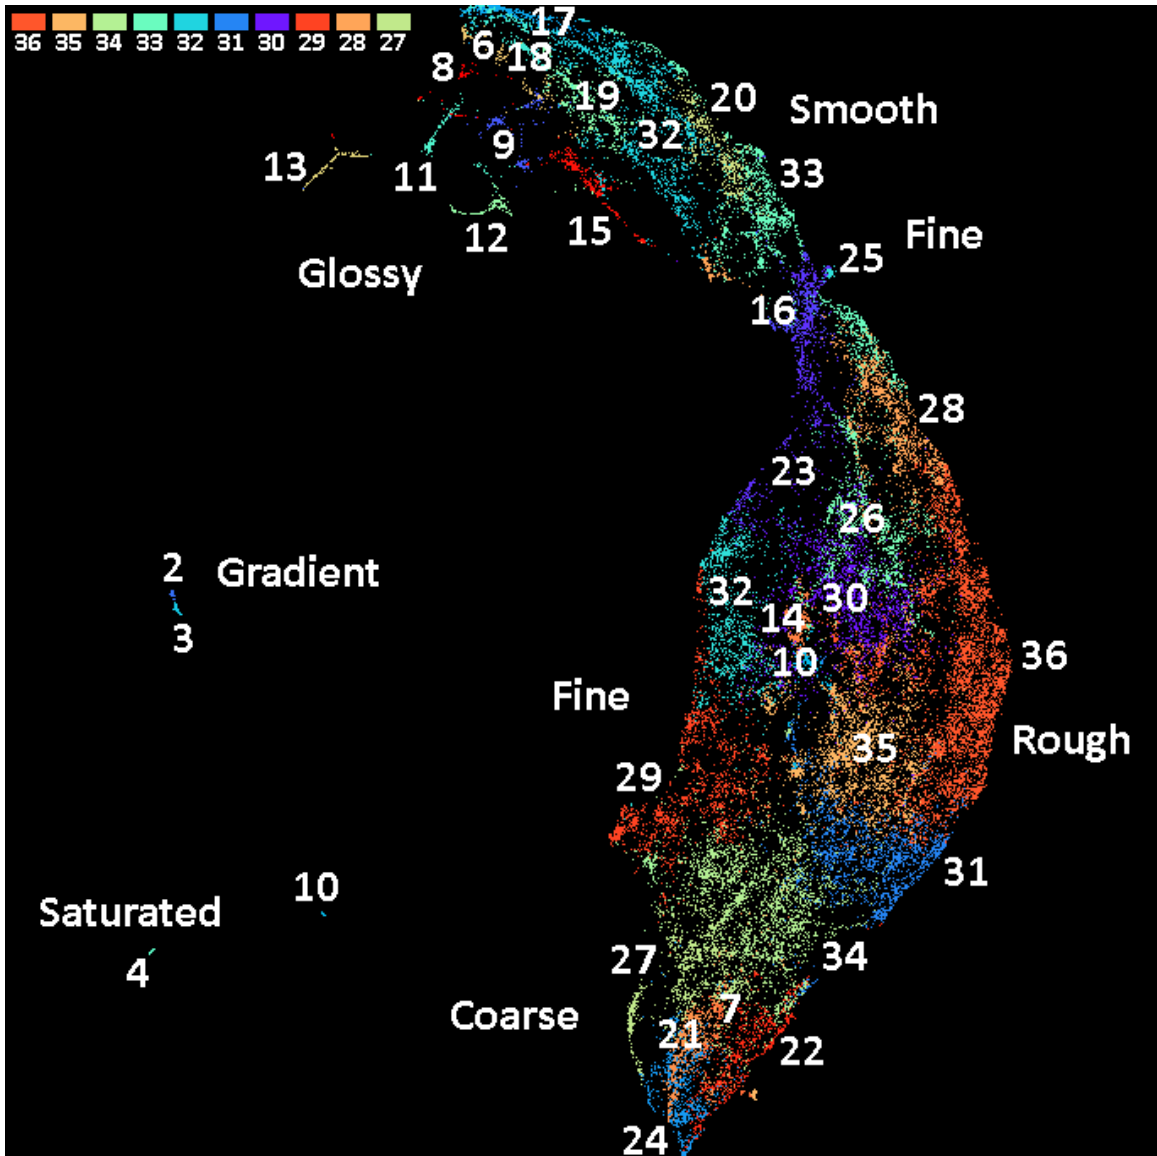



Figure 4.44: UMAP visualization of the Recoat Post-Melt Kylberg-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.

We observe a sparser projection with similar topology in that the same global transition from coarse to glossy surfaces is preserved. Unlike in the Kylberg-TCNN experiment, bulky regions break into more equally-sized clusters that isolate transitions and subcategories with a higher textural homogeneity. For example, cluster 36 contains rough surfaces, but these become more uniformly coarse when transitioning through clusters 31, 34, and 22. We see a similar behavior for smooth to glossy surfaces through clusters 33, 20, 32,


and 8. Notice that cluster 32 is not contiguous in the projection and is split among smooth and fine surfaces. We suspect the smooth and fine categories may be closer neighbors in feature space than the projection is conveying.

For consistency, we also look at the effect of coupon positioning on melted surfaces by coloring the projection by coupon row, column, and index in Figure 4.45.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.45: UMAP visualization of the Recoat Post-Melt Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

Coupon row separation is clearer in this projection while coupon column remains non-influential. The majority of the separation does not separate coupons between surface categories but among surface categories with the exception of finer surfaces primarily belonging to middle row coupons 4, 5, and 6. As we transition to the ALOT experiments, we evaluate whether the homogeneity and utility of their clusters match these.

ALOT-TCNN

Proceeding with the ALOT-TCNN, we cluster the recoat post-melt feature space with a noise threshold of 0 yielding 44 clusters, the most out of all experiments. We relabel these clusters sorted by size and report them in Figure 4.46.






Figure 4.46: Recoat post-melt cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.


This distribution is similar to the one in the Kylberg-WCNN experiment in that there are many large similar-sized clusters; however we do have 3 significantly larger clusters with less textural homogeneity. Most clusters 1 to 34 are small with some likely being too specific. We inspect randomly sampled images from each cluster in Figure 4.47 to understand why the distribution is different.




(1) Glossy etching; shiny top edge




(2) Coarse; mostly bright




(3) Coarse and wavy




(4) Glossy; raised corner-
s/edges; specks




(5) Smooth etching; bright cor-
ners/edges; specks




(6) Glossy etching; bright edges




(7) Rough; overall bright; dark
edges




(8) Fine bumpy




(9) Smooth; bright spot-
s/edges/corners




(10) Coarsely rough; bright spots (11) Glossy etched; raised corners/edges (12) Glossy etched; raised corners




(13) Solid; dark spots (14) Flat etching; small bright corners/edges (15) Glossy; raised corners/edges




(16) Shiny; large gradients (17) Coarse; dark reliefs; specks (18) Glossy; shiny corner gradients




(19) Grooved; bright corners/edges




(20) Semi-rough; specks




(21) Glossy; raise corners/edges




(22) Uniformly coarse




(23) Grooved rough; dark edges bright




(24) Coarse and wavy; mostly




(25) Fine; rough specks




(26) Smooth; rough specks




(27) Fine etching; light specks




(28) Flat; rough specks




(29) Grooved rough




(30) Glossy; bright edges/-corners; specks




(31) Coarse; mostly bright




(32) Flat; rough specks




(33) Fine etching; bright spots



(34) Semi-rough



(35) Grooved and etched; bright rough spots



(36) Glossy and semi-rough; specks

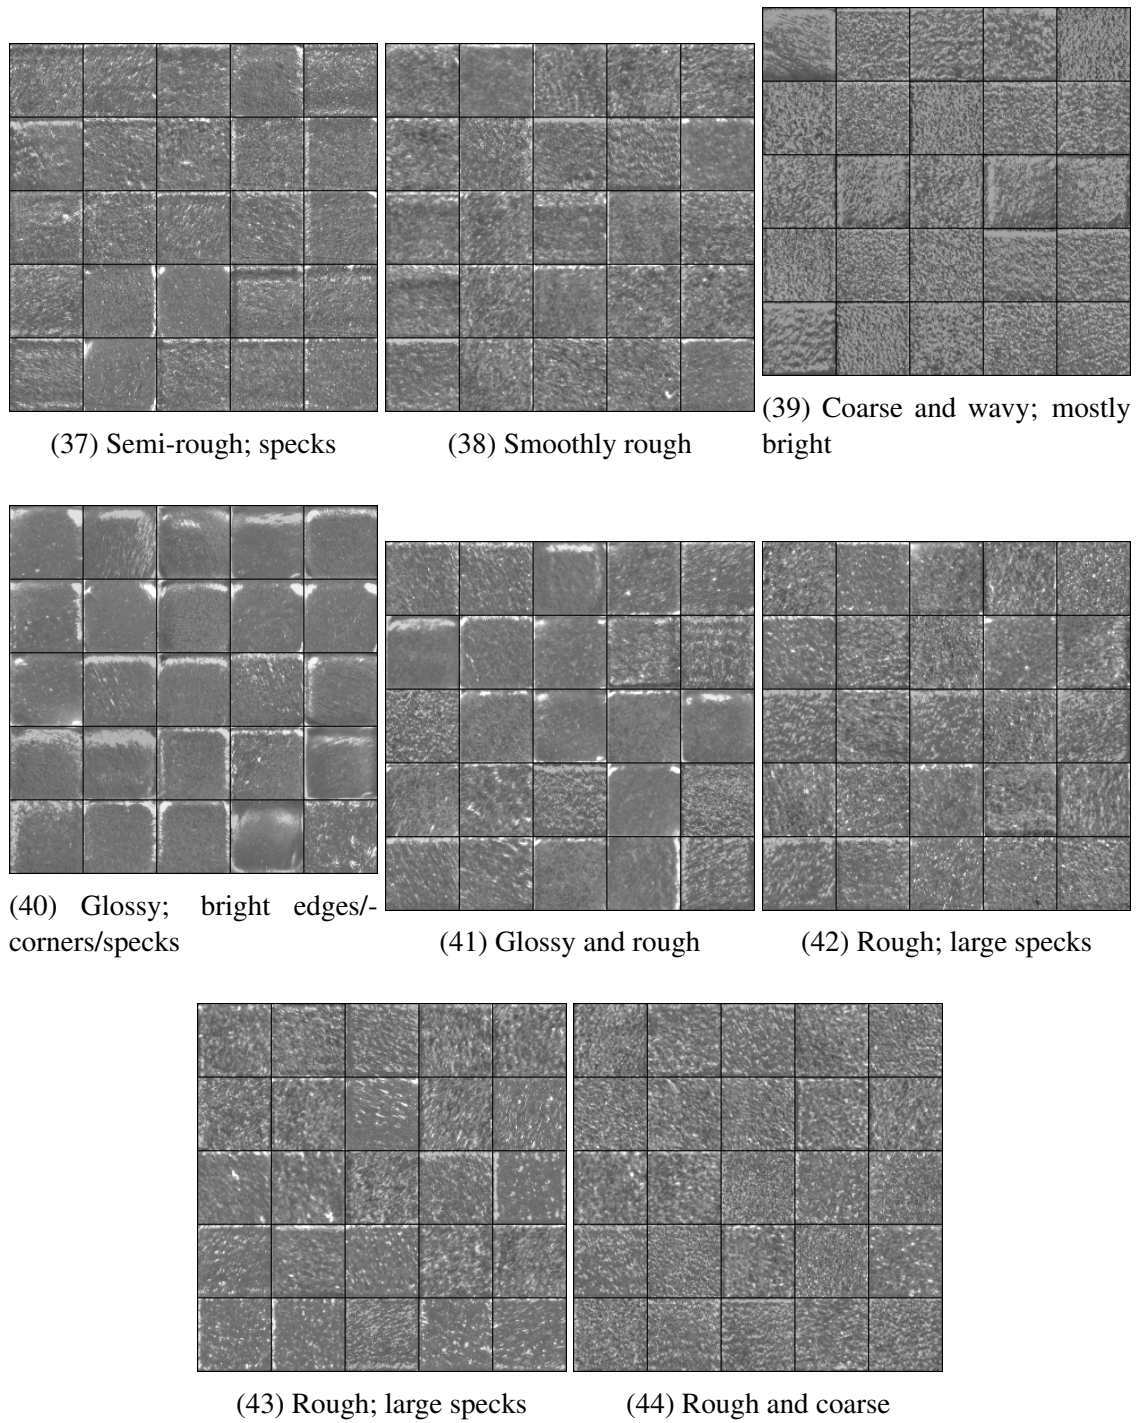


Figure 4.47: Normalized images sampled from each cluster in the recoat post-melt ALOT-TCNN feature space.

We see many similar clusters from to the Kylberg-WCNN experiment, with a few unique additions. Cluster 1 has a mere 20 samples that are nearly identical. This cluster is

distinct, but would be better grouped with cluster 6 since samples from both clusters capture etched glossy surfaces. Cluster 17 introduces a new useful subcategory of coarse surfaces with dark reliefs. Unlike in the Kylberg-WCNN experiment, cluster 16 containing shiny surfaces no longer splits unnecessarily. The largest clusters 42 to 44 appear reasonably homogeneous in the rough category, but clusters 35, 36, and 41 are questionably intermingling smoother and rougher surfaces. We see clearer separation for fine surfaces among clusters 27, 28, 32, and 33. This clustering competes with that in the Kylberg-WCNN experiment as it segments finer surfaces cleaner but introduces more textural confusion among smooth and rough surfaces. We look to the UMAP projection in Figure 4.48 colored by cluster to see how these observations are reflected spatially.

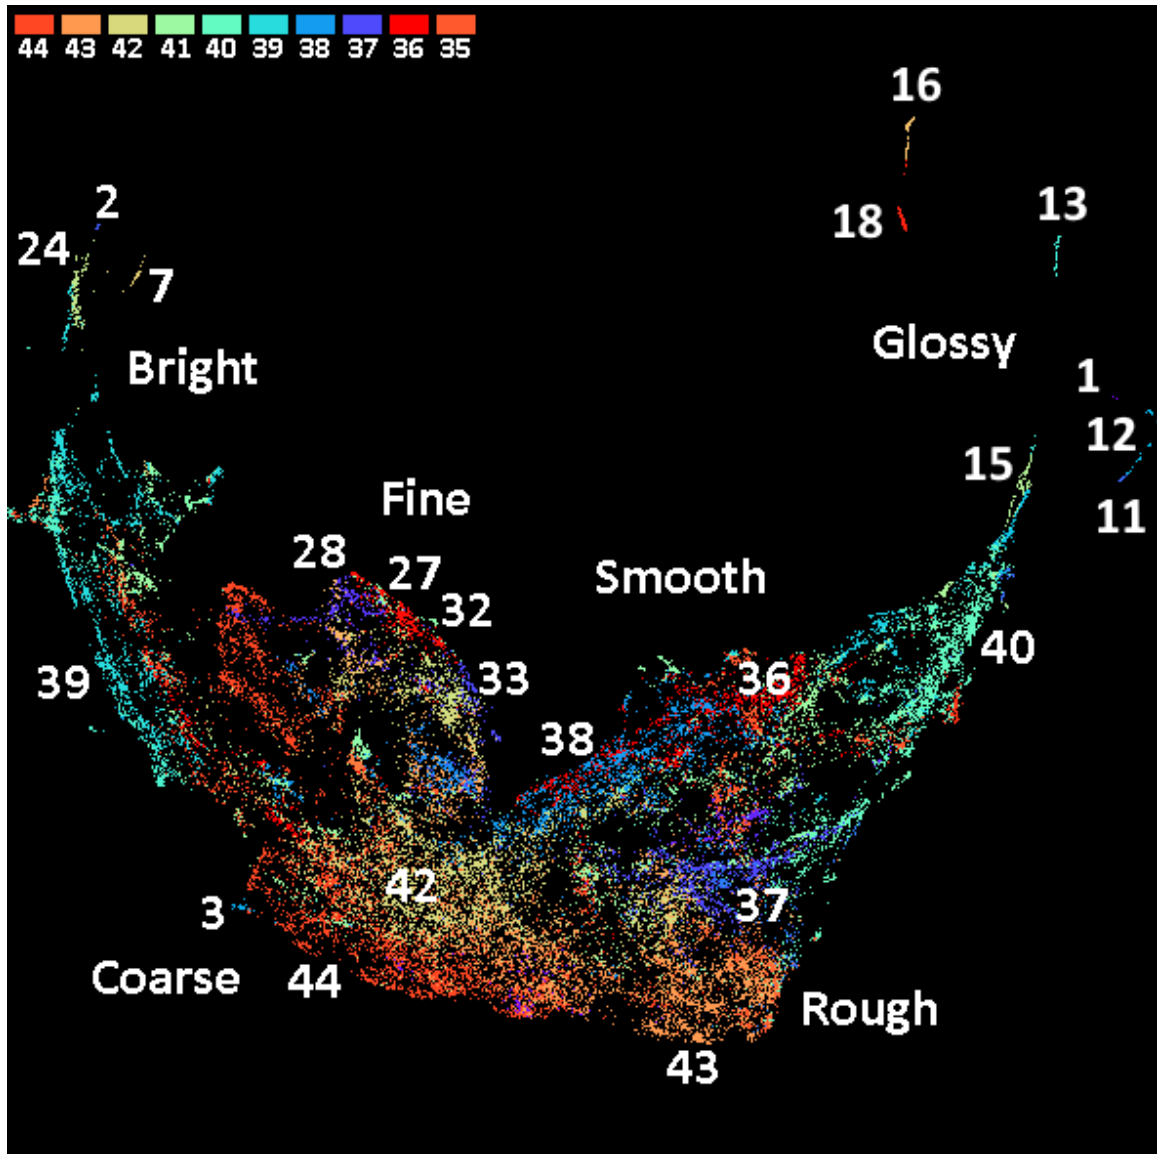



Figure 4.48: UMAP visualization of the Recoat Post-Melt ALOT-TCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.

We observe a significantly different shape compared to the Kylberg experiments; however the topology appears to be consistent as it probably is data-driven and varies primarily at a local level depending on the textural representation. We confirm a few initial observations. Clusters 35, 36, and 41 are not contiguous and spread across the smooth, rough, and coarse regions in the projection, which we saw was reflected in the clustering. The fine surface region pulls away from the rest of the projection and segments cleaner over the


composition of clusters 25, 27, 28, 32, and 33.

The projection is overall sparser. We can see some spatial separation between the major surface categories such as coarse from rough. Many clusters are isolated away from the main group with the glossy surfaces being clusters 13, 16, and 18, and the coarse surfaces being clusters 2, 7, and 24. Note that 7 is not as similar to cluster 2 and 24 for being coarse than it is for being overall bright. Part of the transition through the coarse region leads to this over-saturated coarse surface subcategory. We finish our discussion for this experiment by examining the effect of coupon positioning by coloring this projection by coupon number in Figure 4.49.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.49: UMAP visualization of the Recoat Post-Melt ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

This coloring, like the previous ones from the Kylberg experiments, reveals the influence of build plane row positioning. In the Kylberg projections, the coarse region tail was two-sided with one side containing saturated textures, but in this projection the non-saturated samples do not extend alongside the brighter half. The saturated samples belong to the front row coupons, while the dimmer samples belong to the back row coupons. This

is why coupons 1, 2, and 3 do not extend far into the fine or coarse surface regions in this projection. It is suspected the back row coupons 1, 2, and 3 had insufficient lighting reflections to produce this behavior within that region of the build plane. Lastly, we cross validate from the Kylberg experiments that middle-row coupons 4, 5, and 6 primarily make up the fine surface region. This concludes our analysis for this experiment directing our attention to the ALOT-WCNN texture representation.

ALOT-WCNN

Finishing with ALOT-WCNN, we cluster the recoat post-melt feature space using a noise threshold of 0 yielding 25 clusters. We relabel these clusters sorted by size and report them in Figure 4.50.

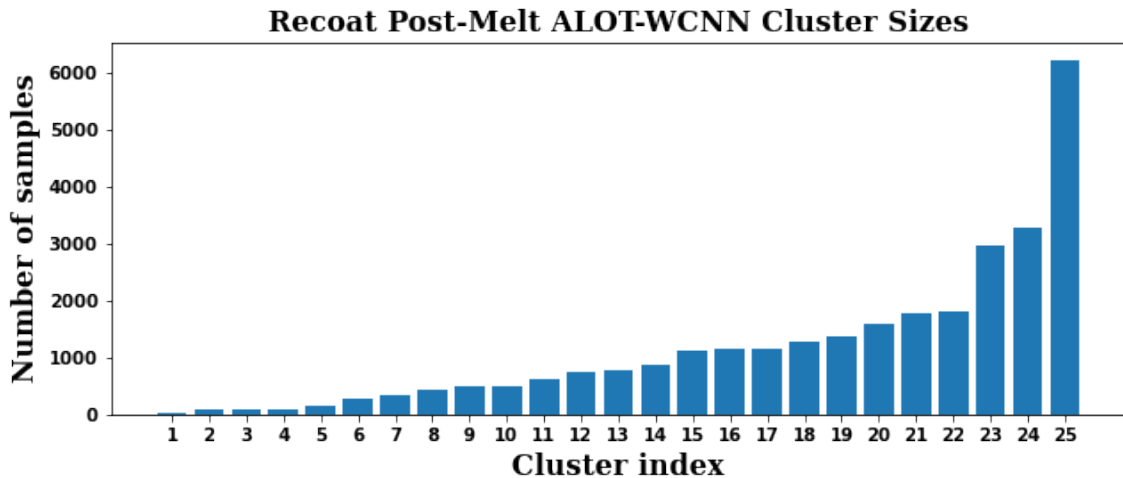




Figure 4.50: Recoat post-melt cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.

This distribution has a nearly identical shape to the Kylberg-WCNN, but the clusters are larger in general. We expect the quality of this clustering to be somewhere between the Kylberg-TCNN and Kylberg-WCNN. In Figure 4.51, we randomly sample and label images from each cluster.




(1) Glossy etching; raise cor-
ners


(2) Glossy etching; raised cor-
ners




(3) Solid; dark spots




(4) Grooved; dark edges




(5) Shiny; large gradients




(6) Wavy; mostly bright




(7) Coarse; mostly bright



(8) Fine etching; specks




(9) Semi-rough



(10) Rough; specks

(11) Coarse and wavy

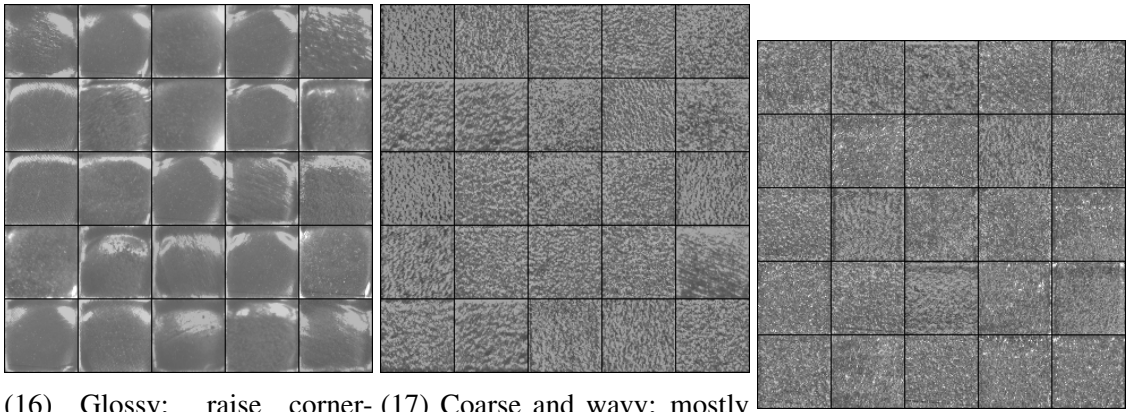
(12) Smooth; rough specks



(13) Fine; rough specks

(14) Glossy; bright corner-
s/edges

(15) Coarse; rough specks



(16) Glossy; raise corner-
s/edges

(17) Coarse and wavy; mostly
bright

(18) Fine; bumpy

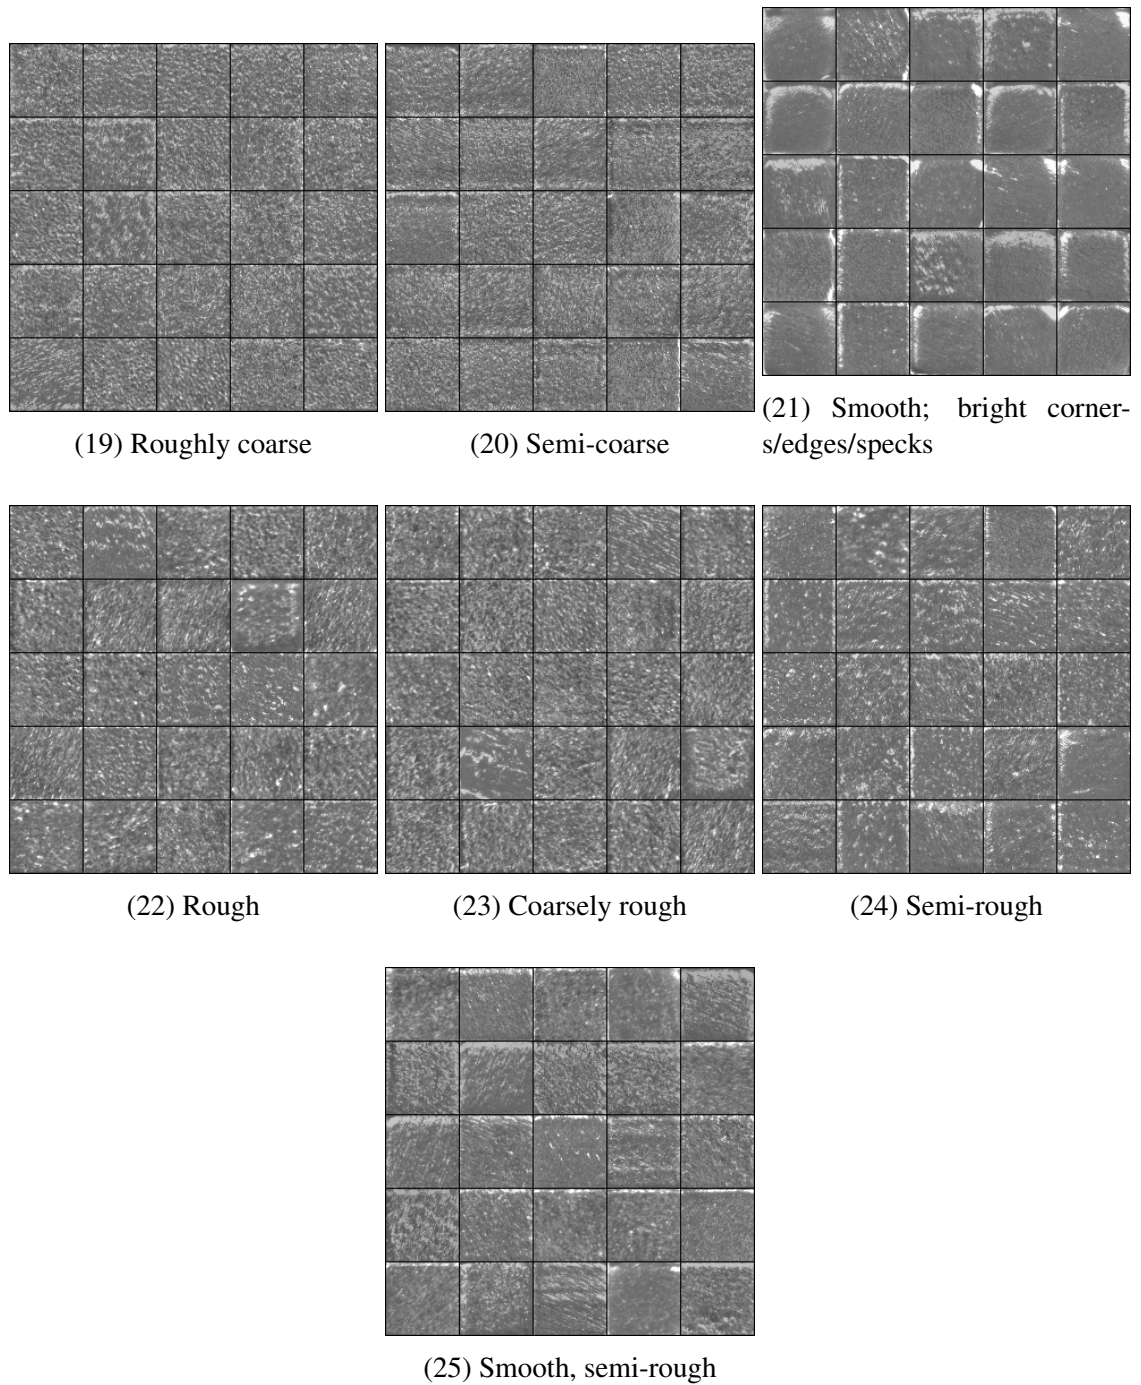


Figure 4.51: Normalized images sampled from each cluster in the recoat post-melt ALOT-WCNN feature space.

We observe many similarities to previous experiments. In the ALOT-TCNN clustering, smooth and rough classes were being mixed together in many clusters. A less severe case occurs in the largest cluster 25 whose samples are a mix between those in clusters

12 and 24. In the Kylberg-WCNN clustering, coarse and rough subcategories were established, and those persist among clusters 7, 10, 11, 15, 17, 19, 20, 22, and 23. Across all experiments, glossy and shiny surfaces isolate well, especially those in clusters 1, 2, 3, 5, 14, and 16. This is also the case for fine surfaces in clusters 8, 13, and 18. These clusters demonstrate sufficient segmentation of surface categories. We verify the spatial relationships of these clusters are still similar to those in previous experiments by coloring a UMAP projection of this feature space by cluster in Figure 4.52.





Figure 4.52: UMAP visualization of the Recoat Post-Melt ALOT-WCNN feature space. Minimum distance of 0.00 and 50 neighbors were used.


Once again, the topology is similar in that there is a global transition from coarse to glossy surfaces. The overall shape is quite different simply due to stretching and warping. This elbow shape is realized by the close neighboring of the fine and smooth surface regions causing the entire projection to bend.

We observe a few transitions. The transition within the rough region from coarse to smooth is realized by the chain of clusters 23, 22, and 24. Cluster 25 shortcuts this transition with less textural homogeneity. Smooth surfaces transition into glossy surfaces from clusters 12 and 21 to 14 and 16 with the special varieties in clusters 2, 3, and 5. Fine surfaces transition into coarse surfaces from clusters 8 and 20 to 15, 18, and 17. This concludes our analysis of these clusters. For completeness, we show the projection colored by coupon in [Figure 4.53](#). We observe coupon positioning has a similar influence as mentioned from the ALOT-TCNN experiment.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.53: UMAP visualization of the Recoat Post-Melt ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

Summary

In the recoat post-melt data, we effectively clustered broad categories of surfaces ranging from glossy, smooth, fine, rough, and coarse. Among these, we saw additional separation at transition areas within the projections and subcategories exposing over-saturation, etc-

ing, and grooves. Visualizing the spatial relationships of clusters obtained by FDC using UMAP projection verified both algorithms agreed on the neighboring of surface categories. Coupon positioning had less influence on the recoat post-melt feature space structure compared to the recoat post-spread feature space. We observed coupon images grouped together row-wise but not individually. Fortunately, this grouping was primarily within surface categories. If the group was across surface categories, the position on the build plane would be the leading factor for clustering surfaces which would nullify any attempt of identifying the true physical surface roughness of the melted material using the same machine setup. There was one exception in the clustering. The finer surfaces category was primarily composed of middle row coupons likely due to lighting and camera focus.

From these results, we recommend similar future modifications as mentioned in the recoat post-spread analysis. Lighting still influenced the ability to accurately isolate surfaces, but dead pixel residue from sensor damage was negligible, thus not a priority. We experienced a new obstacle, namely over-saturation, which led to surfaces with neighboring each other in projections because they shared over-saturation features despite some of their other dissimilar surface features. Lowering the light level may resolve this issue.

Among these experiments, smaller clusters containing more specific instances of surface categories were nearly identical. Larger clusters primarily contained rougher surfaces which were best split by the Kylberg-WCNN texture representation. Although the ALOT-TCNN produced the most clusters and isolated a unique subcategory of coarse surfaces, there was some confusion between smooth and rough surfaces within larger clusters. The distribution of raised points and bright specks on smooth surfaces contributed to this association with rough surfaces. We established the Kylberg-TCNN clustering as a baseline and reserve the Kylberg-WCNN clustering for further research.

In summary, the recoat post-melt data was cleaner and easier to achieve our goal of isolating general categories of melted surfaces. The following modalities do not use the same sensor nor operate in the same optical spectrum; therefore we do not expect lighting to

influence their clusterings. In the next section we analyze the thermal tomography data.

4.2.3 Thermal Tomography

The thermal tomography data consists of thermograms from each coupon on each layer. These thermograms are known to reveal spatter, cold spots, and roughness in the energy emissions. Spatter are hot ejected particles of powder that release a lot of energy. They are manifested as bright blobs in thermal imagery. Cold spots are low-energy density areas. They may imply there is a lack of fusion in the powder. Some instances are induced by severe spatter on previous layers. They appear as dark reliefs in thermograms. This modality can experience texture washout when the entire region of interest emits too much energy causing over-saturation. This typically happens when the energy density of the processing parameters is too high. A similar scenario occurs when the energy density is too low. In both cases, they yield no texture as all the pixels are the same value. We exclude these instances from our analysis.

In this section, we perform clustering to group similar thermograms. This clustering would be beneficial if thermograms segment by spatter, cold spots, and levels of roughness. We begin by assessing how many clusters the thermal tomography feature space can break into while varying the noise threshold of FDC. We display the number of clusters per threshold for each experiment in [Figure 4.54](#).




Figure 4.54: Plot of the average number of clusters the thermal tomography data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.

At most noise thresholds, the Kylberg experiments consistently produce more clusters than the ALOT experiments. We hypothesize the Kylberg clusterings will be finer with a higher textural homogeneity in general. In following experiment sections, we evaluate whether this is the case and whether any of the ALOT experiments have an advantage despite producing fewer clusters.

Kylberg-TCNN

Starting with Kylberg-TCNN, we cluster the thermal tomography feature space using a noise threshold of 0 which yields 40 clusters, the most out of all experiments. We relabel these clusters sorted by size and report them in Figure 4.55.






Figure 4.55: Thermal tomography cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.


This distribution reveals several clusters 1 to 19, 20 to 30, and 31 to 36 are relatively uniform in size to each other. We expect the smaller of these clusters to capture less common phenomena with higher textural homogeneity. Clusters 37 to 40 are noticeably larger than other clusters. We expect them to contain more common and broader categories. We inspect these clusters by randomly sampling and labeling their images in Figure 4.56.




(1) Coarse




(2) Semi-rough, edge gradients




(3) Solid bright border; large dark region




(4) Solid bright border; streaked dark region




(5) Solid border; dark regions




(6) Smooth grains; bright corners; dark edge




(7) Fine; subtle gradient




(8) Smooth; grains; bright corners; gradients




(9) Solid; small dark regions




(10) Solid; small dark regions




(11) Smooth; dark grains; bright spots/corners




(12) Fine, semi-coarse; edge gradients




(13) Smooth; large dark edge/- corner




(14) Smooth grains; bright corners




(15) Smooth; grains; gradients; bright corner




(16) Smooth; large gradients; bright corner




(17) Solid; dark specks




(18) Smooth; large gradients; bright corner




(19) Finer coarse; dark reliefs




(20) Smooth; subtle grains; bright spots




(21) Smooth; rough grains; bright corners




(22) Solid corners; dark grained relief




(23) Coarse; bright/dark spots




(24) Smoother semi-rough; edge gradients




(25) Semi-rough; edge gradients




(26) Semi-rough; edge gradients




(27) Smooth; subtle grains; bright corners




(28) Smooth grains; bright spots/corners




(29) Semi-rough; grains; spots




(30) Fine, semi-coarse




(31) Fine; reliefs




(32) Smooth grains; bright spot/corner




(33) Finer coarse



(34) Smooth; edge gradients



(35) Fine; gradient reliefs; dark edge



(36) Rough; bright/dark spots

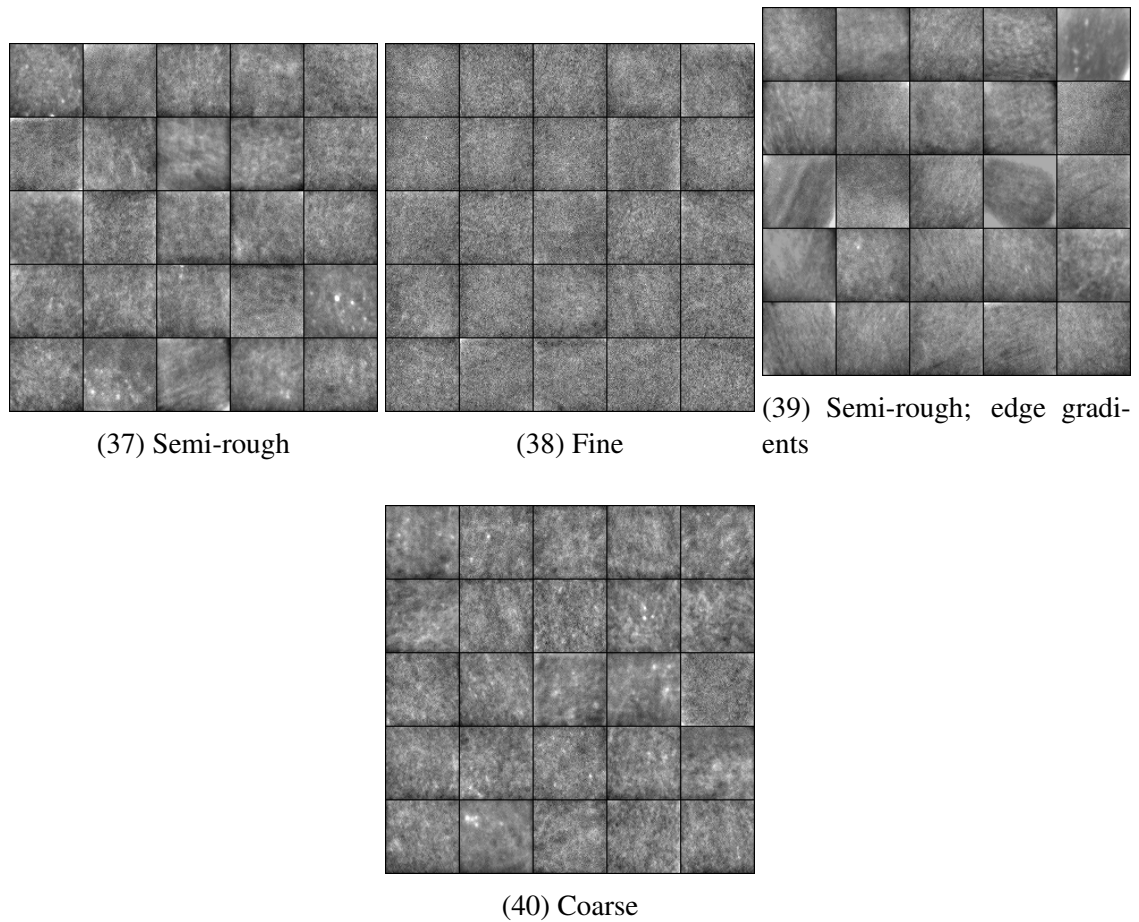


Figure 4.56: Normalized images sampled from each cluster in the thermal tomography Kylberg-TCNN feature space.

By inspection, most clusters are texturally homogeneous regardless of size. They primarily isolate textures by levels of roughness. We name many clusters corresponding to these general texture groups. Finer samples are observed in clusters 7, 12, 19, 30, 31, 35, and 38. Coarser samples are distributed among clusters 1, 12, 19, 23, 30, 33, and 40. Rougher samples are grouped into clusters 24, 25, 26, 29, 36, 37, and 39. Smoother samples span clusters 6, 8, 11, 13, 14, 15, 16, 18, 20, 21, 27, 32, and 34. Finally, low frequency textures with solid regions are spread across clusters 3, 4, 5, 9, 10, 17, and 22. In Table 4.4, we summarize which clusters capture different roughness levels and process phenomena. From this table, we note there exists a lot of textural overlap such that many clusters could be merged. This hard labeling is not very accurate, but a fuzzy labeling could be.

Texture	Clusters
Fine	7, 12, 19, 30, 31, 35, 38
Coarse	1, 12, 19, 23, 30, 33, 40
Rough	36
Semi-rough	2, 24, 25, 26, 29, 37, 39
Smooth	6, 8, 11, 14, 16, 18, 20, 21, 22, 24, 27, 28, 32, 34
Saturated	3, 4, 5, 9, 10, 17, 22, 27
Grains	4, 6, 8, 11, 14, 15, 21, 22, 27, 28, 29, 32
Gradients	2, 6, 8, 12, 13, 15, 16, 18, 19, 24, 34, 35, 39
Center-relief	3, 4, 5, 9, 10, 22, 27
Spatter	2, 8, 11, 20, 28, 29, 32, 36, 37, 40

Table 4.4: Table of process phenomenon seen in Kylberg-TCNN thermal tomography clusters.

Regarding process phenomena, we acknowledge spatter is not isolating well. In many clusters where some spatter is present, we suspect the background pixels are the driving force for segmentation. For example, compare the smooth cluster 20 and the rough cluster 36. While both clusters contain some samples capturing spatter, the overall roughness of those samples matches the roughness category of the other samples in their respective cluster. Cold spots were lacking in this dataset, but the clustering can segment samples by the depth of reliefs. We observe this when comparing cluster 19 and cluster 38. To avoid confusion, we clarify samples with a large central dark region such as those in clusters 3 and 4 are not considered cold spots. This is a separate phenomenon created by powder being bulldozed by the laser to the extent that its height obstructs the camera's view of the weld emissions.

In general, this clustering adequately distinguished thermograms by levels of roughness as well as other features such as gradients and bright/dark corners/edges, but it could be better in two ways. Several near-identical clusters should be grouped together such as clusters 5, 9, 10, and 17. Spatter should segment into its own cluster. We proceed analysis by visualizing the spatial relationships between clusters and categories of thermograms depicted from a UMAP projection of this feature space colored by cluster in Figure 4.57.

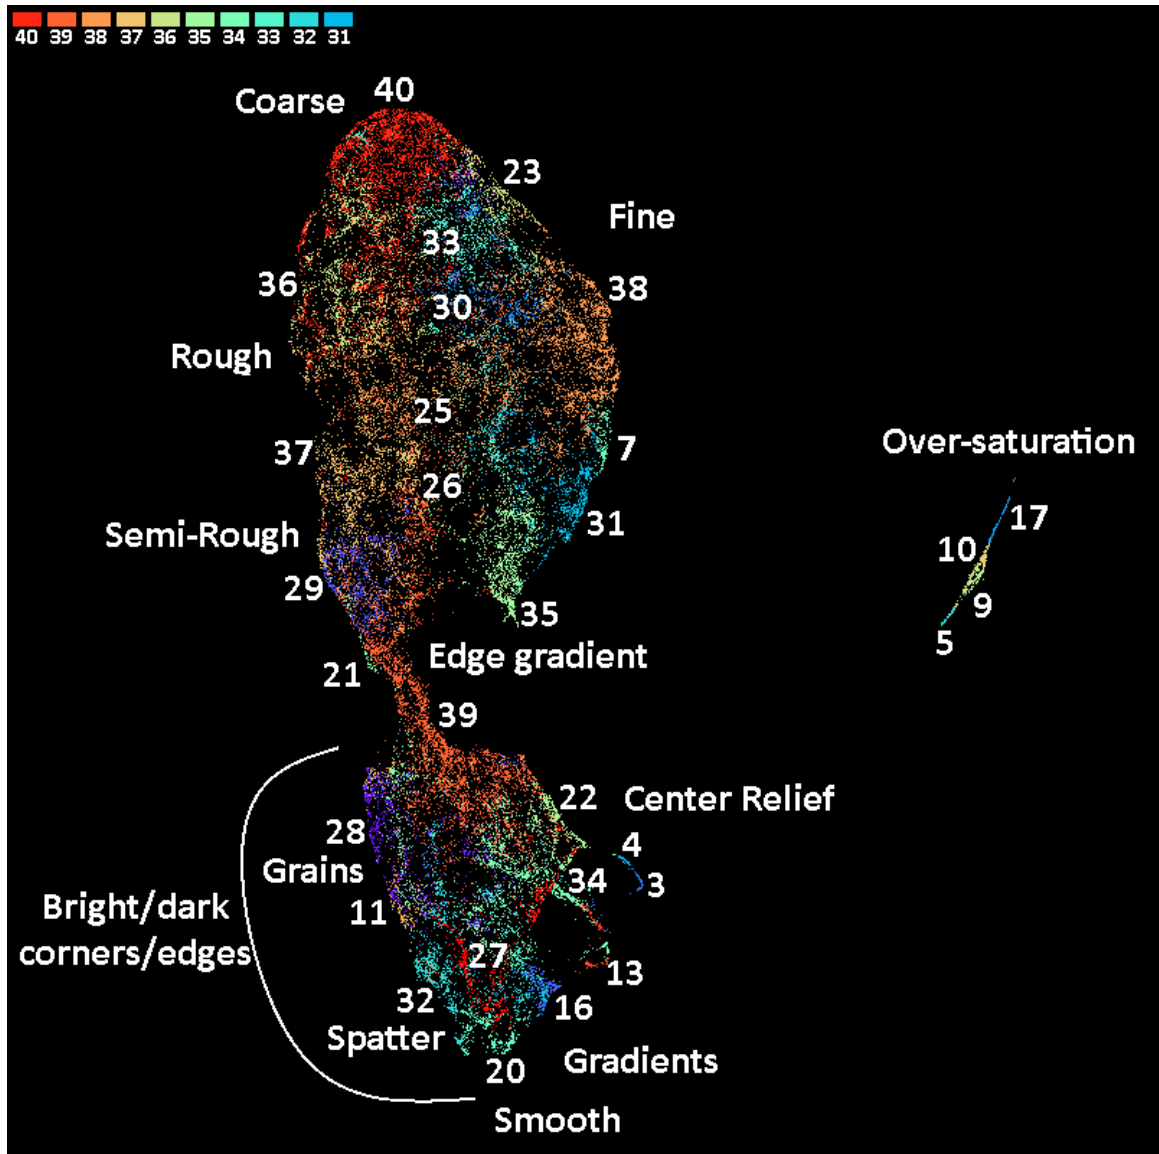



Figure 4.57: UMAP visualization of the Thermal Tomography Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

From previous sensor modalities, we have observed UMAP projects most of the data into a single contiguous cluster. Thermal tomography is no exception, but the appearance of two major clusters bridged by cluster 39 is visible. From the recoat post-melt data, we saw general categories of surfaces grouping together. In this projection, we observe a similar relationship. Fine thermograms transition to coarse ones through clusters 35, 31, 7, 38, 30, 33, 23, and 40. Similarly, samples get rougher and smooth out when passing through


clusters 36, 25, 37, 26, 29, and 21. Most clusters near the bottom region in the projection are smoother with bright or dark corners and edges. Some of these samples feature subtle or deep grains. Others feature gradients and a large central dark relief. Most spatter falls to the bottom of the projection along with smooth samples. This entanglement is observed when looking at the samples from clusters 20 and 32. Over-saturation is more common in this dataset, so we see a larger strand of it in the projection composed of clusters 5, 9, 10, and 17. These clusters vary in the amount of pixels having the same value.

This projection reveals an intuitive topology describing the arrangement of general categories of thermograms. We observe a lot of textural overlap and transitions between categories. This appears to be the case for most datasets. For consistency, we evaluate whether coupon positioning on the build plane influences the clustering by coloring the projection by coupon row, column, and index in [Figure 4.58](#).



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.58: UMAP visualization of the Thermal Tomography Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

This visualization reveals coupon positioning may have an influence in the top cluster, but otherwise coupon samples are equally distributed. The primary observation is coarse samples from coupon 4, 5, and 6 in the middle row are less common. This concludes our analysis for the Kylberg-TCNN experiment. Overall the results were quite pleasing. In the following experiments, we talk about consistencies and differences. We anticipate the

segmentation of spatter in the one of the following experiments may be better.

Kylberg-WCNN

We cluster the Kylberg-WCNN feature space with a noise threshold of 0 which produces 33 clusters. We relabel these clusters sorted by size and report them in Figure 4.59.






Figure 4.59: Thermal tomography cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.


This distribution is similar to the Kylberg-TCNN but differs in that instead of 4 noticeably larger clusters, there is only one cluster 33. We suspect some samples from the larger clusters have spread across the smaller clusters. This clustering will likely be similar. We randomly sample images from each cluster and show them in Figure 4.60.




(1) Smooth; gradients; subtle grains; bright corners (2) Smooth; gradients; subtle grains; bright corners (3) Smooth; gradients; subtle grains; bright corners



(4) Smooth; gradients; subtle grains; bright corners (5) Smooth; gradients; dark corner (6) Fine; edge gradient




(7) Solid; dark specks (8) Smooth bright border; gradients; large dark region (9) Smooth bright border; large dark region



(10) Solid; dark spots

(11) Solid; dark spots


(12) Fine; reliefs



(13) Fine; bright corners; reliefs; gradients

(14) Semi-rough; edge gradients


(15) Smooth; subtle grains; bright corners/spots




(16) Fine; edge reliefs

(17) Smooth; subtle grains; bright corners


(18) Fine; reliefs




(19) Solid bright border;
streaked dark region




(20) Rough, semi-rough




(21) Smooth; edge gradients




(22) Semi-rough; gradients




(23) Smooth; bright corners;
subtle grains




(24) Coarse, rough



(25) Smooth; grains; rough re-
liefs; bright spots



(26) Smooth; edge gradients



(27) Fine, semi-coarse

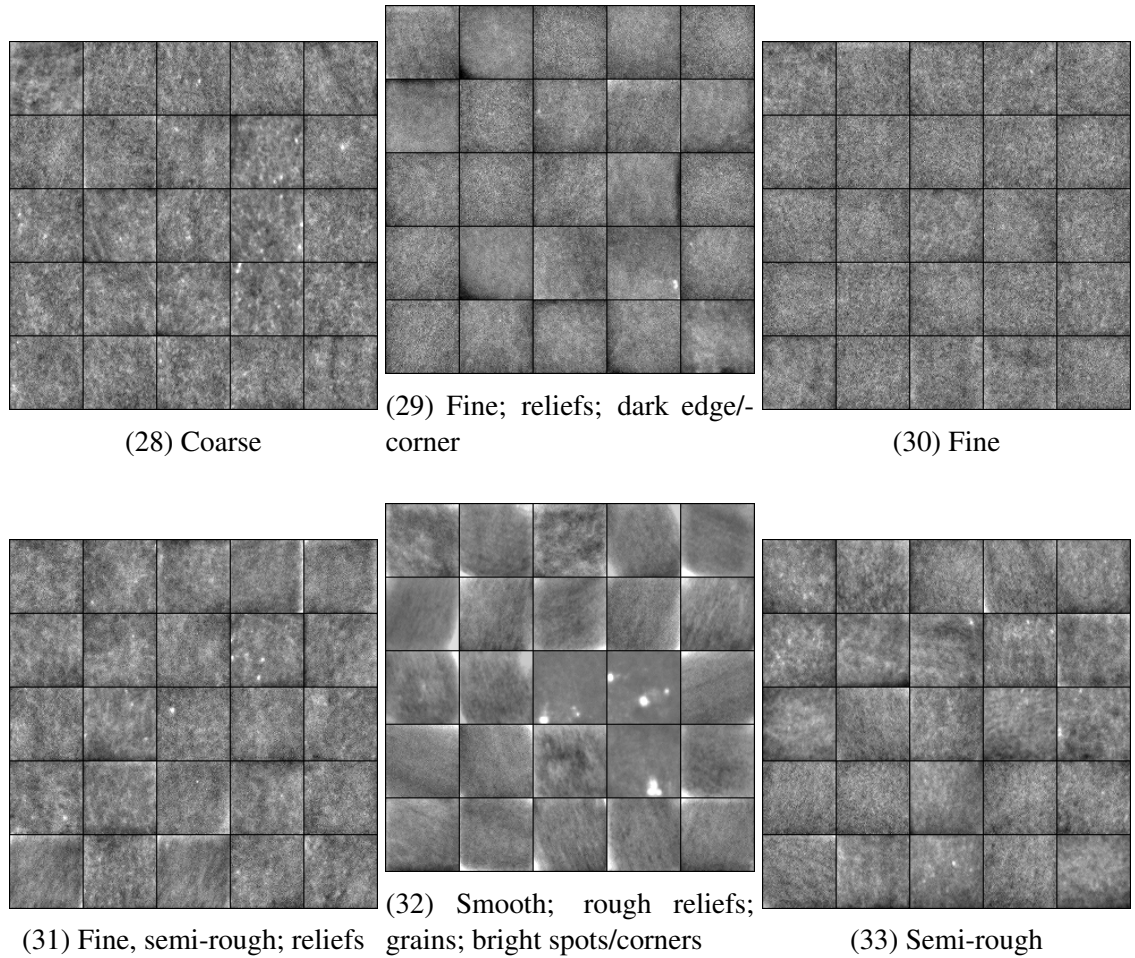


Figure 4.60: Normalized images sampled from each cluster in the thermal tomography Kylberg-WCNN feature space.

We observe similar clusters and categories. For example, cluster 27 is an interpolation between coarse and fine clusters 28 and 30. This scenario is reflected respectively in the Kylberg-TCNN experiment for cluster 23 interpolating between clusters 40 and 38. The presence of gradients and bright or dark corners and edges continues to be a splitting feature. This is observed across fine clusters 29 and 30 and smooth clusters 26 and 32. Solid region clusters 7, 10, and 11 continue to isolate well. In comparison to cluster 13 from the Kylberg-TCNN experiment, cluster 5 isolates similar smooth dark corners and gradients. The textural homogeneity of this clustering is no better than the Kylberg-TCNN experiment. Some clusters could be merged together. An obvious example would be the smallest clusters 1 to 4 as their textural representation is nearly identical. Solid region clusters 7,

10, and 11 could be combined, likewise for fine clusters 16 and 18. Some clusters could be better segmented. Clusters 9 and 23 are primarily smooth but contain over-saturated samples that fit better in cluster 10. For clusters capturing spatter, none purely isolate these anomalies. To see how these clusters are arranged, we color a UMAP projection of this feature space by cluster in Figure 4.61.

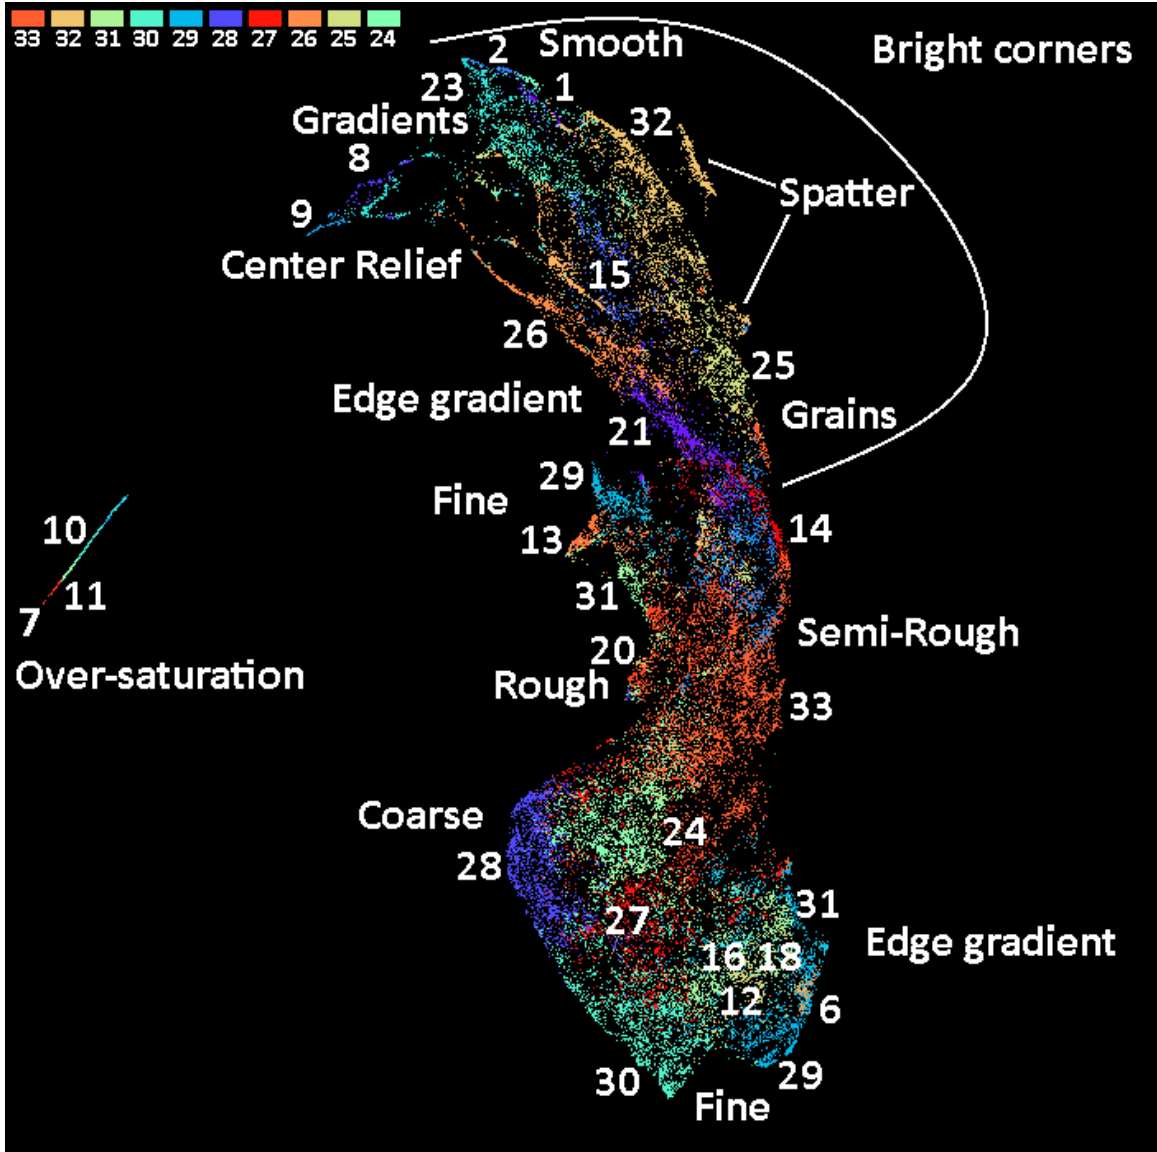



Figure 4.61: UMAP visualization of the Thermal Tomography Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

From this projection, the shape is different with less emphasis on the separation of

two major clusters, but the arrangement of categories is nearly the same. There is a similar transition across the projection from fine to coarse to rough to semi-rough to smooth and finally to over-saturation. There are some differences. The fine region is not contiguous as clusters 29 and 31 stretch above the rough region. This disconnected region primarily captures fine samples with bright/dark corners/edges which explains why it is close to the upper half of the projection that captures larger sized features such as gradients, center relief, and spatter. Most spatter is grouped homogeneously into two regions within this projection. The upper cluster that breaks away primarily captures spatter on a smooth background while the other cluster captures spatter on a rougher background with reliefs. We show samples from these spatter regions in Figure 4.62. Unfortunately, both of these regions are grouped into a larger cluster 32 spanning smooth thermograms with grains; therefore, neither Kylberg experiment purely isolates spatter.




(1) Upper cluster; smooth background

(2) Lower cluster; rougher background


Figure 4.62: Normalized images sampled from the spatter regions in the thermal tomography Kylberg-WCNN feature space UMAP projection.

In the Kylberg-TCNN experiment, we saw some influence of coupon positioning, but to verify that is still the case, we color the projection by coupon row, column, and index in Figure 4.63.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.63: UMAP visualization of the Thermal Tomography Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

It appears coupon segmentation is emphasized more in the lower region of this projection. The coarse region is primarily composed of back row coupons 1, 2, and 3. We saw this in the Kylberg-TCNN projection too. The fine region near the bottom of the projection partially splits front row coupons 7, 8, and 9. Inspecting clusters 30 and 29 lead us to suspect gradients might be the differentiating feature. We require more results from other

experiments before our final evaluation on the influence of coupon positioning.

Overall, both Kylberg experiments produced similar clusters and projections that isolated general categories of textures along with other anomalous features such as bright/dark edges/corners/spots, gradients, and grains. Spatter is still a challenge to segment out via unsupervised clustering. Part of the issue could be tied back to bright corners resembling spatter in a spatially invariant textural representation. Another factor could be that the average activation of texture features washes out anomalies with a small footprint. This could be the case because while bright corners are small, they have a large impact on normalization to squash down other pixel values hence making the texture overall smoother. This explains why smoother samples are within the bright corners region. With these concluding statements, we proceed onto the ALOT experiments in the next sections.

ALOT-TCNN

We cluster the ALOT-TCNN thermal tomography feature space using a noise threshold of 0 yielding 22 clusters. We relabel these clusters sorted by size and plot their sizes in Figure 4.64





Figure 4.64: Thermal tomography cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.


This distribution follows an exponential curve. Most clusters are small with a few larger ones, namely clusters 15 to 22. Many subcategories have likely be grouped into larger clusters. We suspect clusters 20 to 22 will have less textural homogeneity because of their size. In Figure [4.65](#) we randomly sample and label images from each cluster to get a visual assessment.



(1) Solid; dark region

(2) Fine; reliefs; bright corner


(3) Smooth; bright spots



(4) Smooth; large dark relief

(5) Smooth; large gradients


(6) Smooth bright border; dark relief




(7) Smooth, coarse; bright spots

(8) Smooth; grains; bright corners


(9) Fine




(10) Fine; reliefs; bright corner/edge




(11) Semi-rough; bright corners




(12) Fine; dark reliefs




(13) Solid; dark regions




(14) Fine




(15) Fine; reliefs



(16) Smooth; grains; spots; bright corners



(17) Fine; reliefs



(18) Smooth, semi-rough

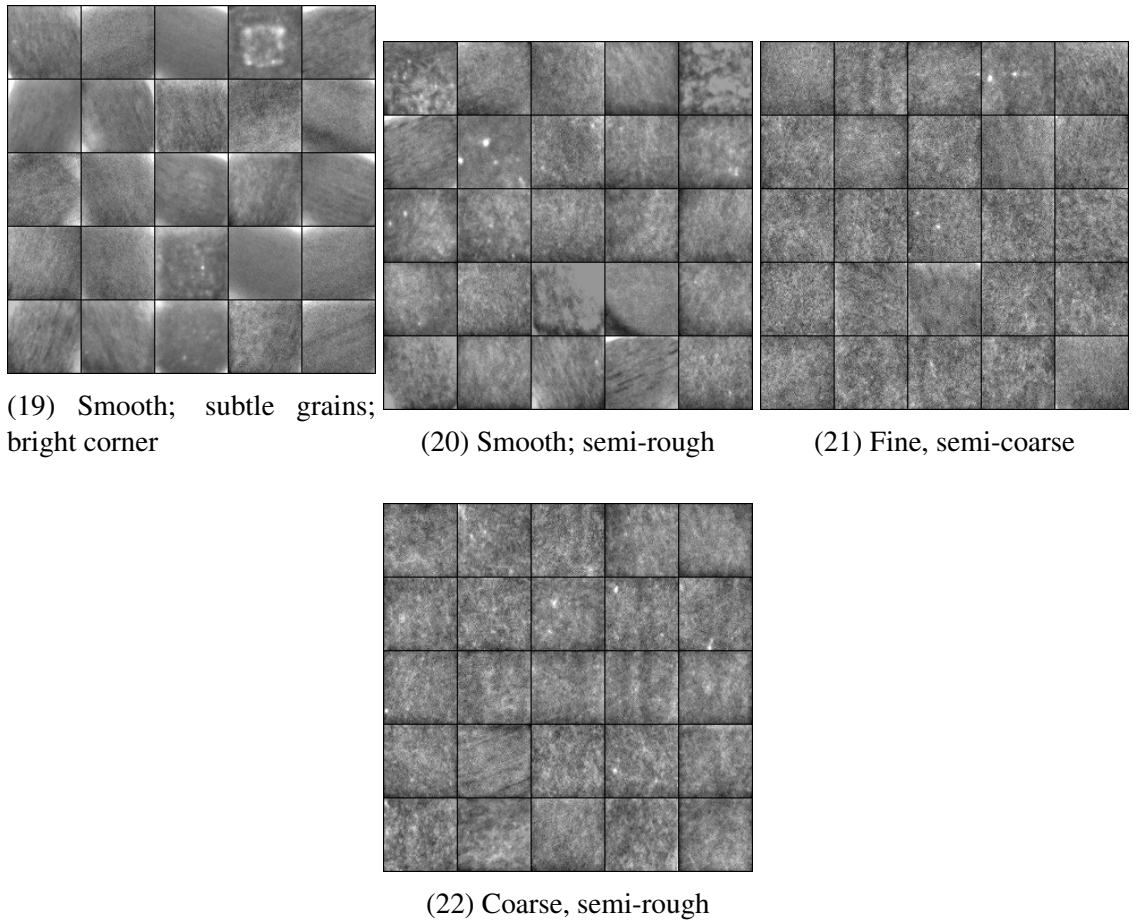


Figure 4.65: Normalized images sampled from each cluster in the thermal tomography ALOT-TCNN feature space.

These clusters capture previous categories; however, there are some differences. We obtain two new small clusters 3 and 7 purely capturing spatter. They overlap differing in the overall roughness. Cluster 7 captures more coarse samples. While these two clusters do not cover all instances of spatter, we can accurately isolate some of it. The cluster containing large dark corners is missing. We saw it as cluster 13 from the Kylberg-TCNN and cluster 5 from the Kylberg-WCNN experiments. The solid category only splits into two clusters, namely cluster 1 for larger dark spots and cluster 13 for smaller dark spots. We assess the textural homogeneity of larger clusters. Clusters 18 and 20 capture a mix of smooth and semi-rough samples. This confusion was also observed in the ALOT-TCNN experiment on the recoat post-melt feature space. Cluster 21 acts almost as transition clus-

ter isolating coarser fine samples. Cluster 22 captures coarser and rougher samples. The transition among these categories is no longer distinct in the clustering. We proceed to visualizing the spatial relationship among clusters in the UMAP projection of this feature space colored by cluster in Figure 4.66.

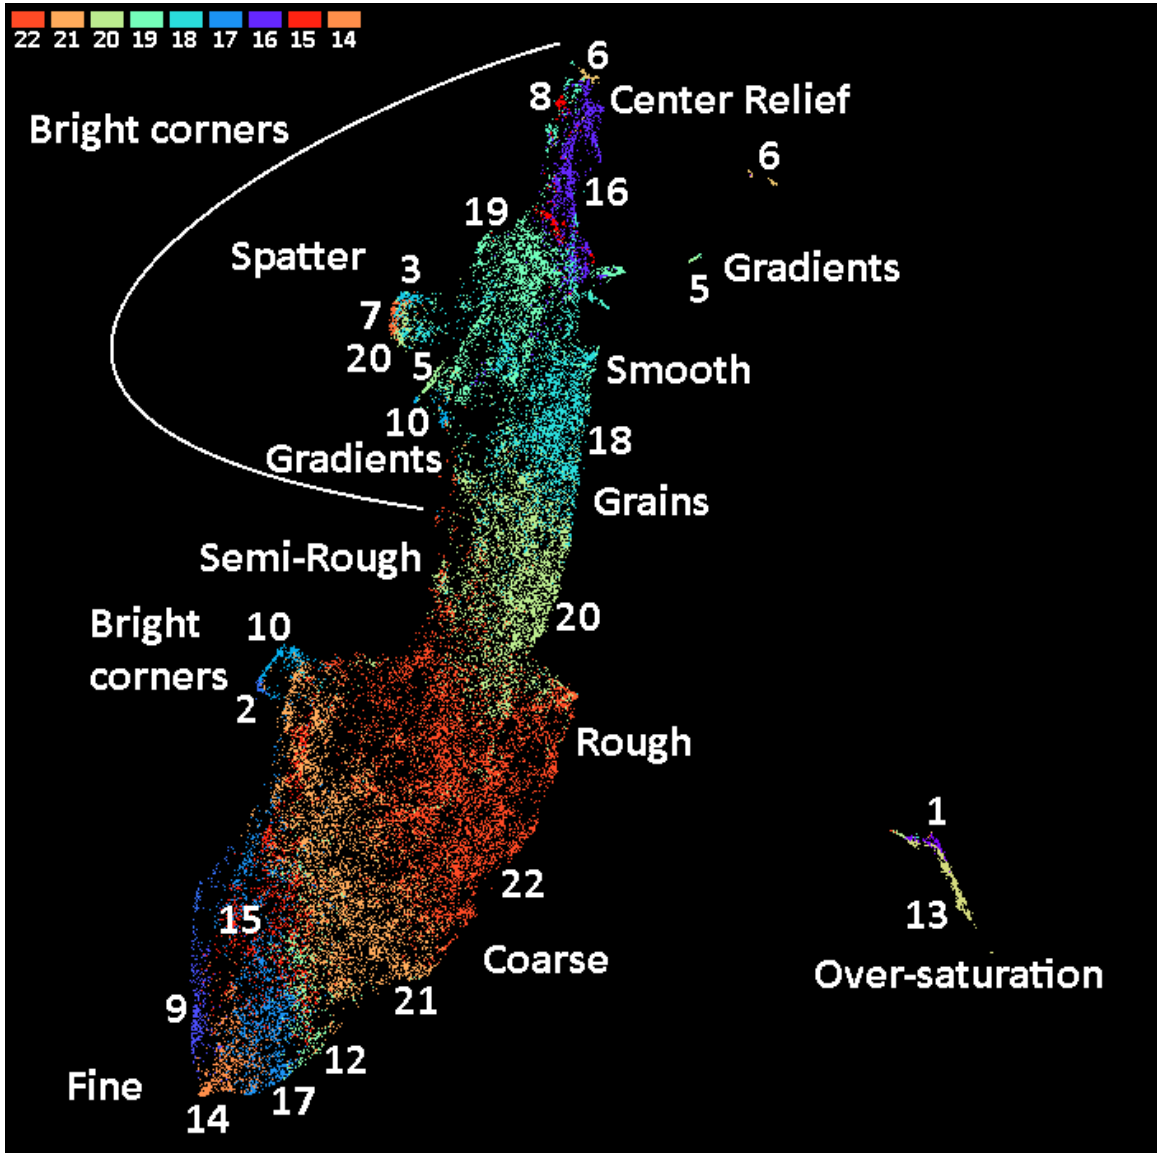


Figure 4.66: UMAP visualization of the Thermal Tomography ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

Like in the Kylberg experiments, we confirm the arrangement of texture categories persists. The fine region wraps around to both the coarse and semi-rough regions. Smooth


regions differ in spatter, gradients, grains, and center reliefs, but all instances have a bright corner or spot. Over-saturation continues to break away. In the recoat post-melt projections, the arrangement of texture categories tended to be consistent across experiments. We expect this to hold for the remaining sensor modalities including this one too.

We discuss clusters next. The spatter region barely breaks off from cluster 18 and 19 containing smooth grained samples with bright corners. It is composed of clusters 3 and 7 as expected by the image sampling, but clusters 18 and 20 are present too. This is because their spatter samples are overall rougher. At the break away part of the spatter cluster, there are samples from clusters 18 and 19 capturing spatter and large dark reliefs. Looking at the larger clusters 21 and 22, they span a wide area of the projection. Cluster 21 resides at the border between fine and coarse. It stretches up towards clusters 2 and 10 where fine samples begin capturing larger features such as bright/dark corners/edges and gradients. Cluster 22 spans most of the coarse and rough regions stretching partially into the semi-rough region primarily dominated by cluster 20 which also extends into the rough region. The transition between coarse and rough categories has been lost in terms of being its own cluster. This is also the case between fine and coarse as well as rough and semi-rough.

A few clusters are not contiguous in the projection. We observe this behavior for clusters 5, 6, 10, and 20. Cluster 5 splits darker more gradual gradients on the left and lighter sharper gradients on the right. The reason the right cluster breaks away from the projection is because its samples do not feature bright corners. Cluster 6 splits barely as samples vary slightly in the presence of grains, but the distance between its clusters is likely a meaningless artifact from UMAP using local distances. Cluster 10 splits finer samples with reliefs and bright corners near cluster 2 and rougher fine samples with grain-like reliefs and bright edges near cluster 5. Some of cluster 10 appears to be taking samples away from cluster 5. Cluster 20 splits because some spatter samples have the same overall roughness as its average sample.

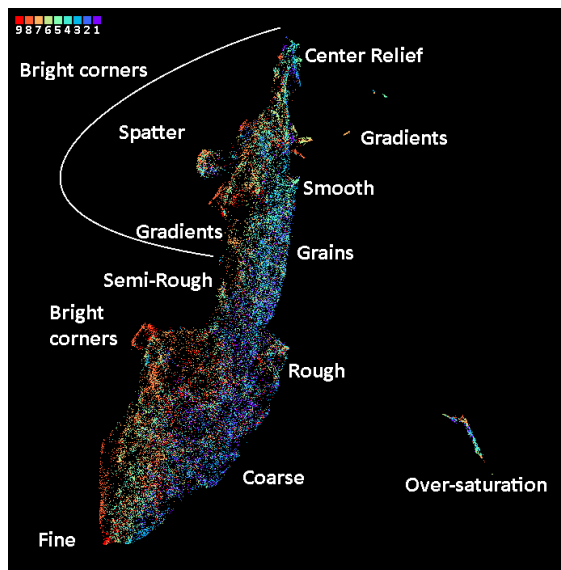
Overall, this clustering does not deviate far from the those in Kylberg experiments. With

less clusters, transitions between roughness categories were not distinguished as their own clusters. The distinction between some smooth clusters were lost. The redeeming factor of this experiment is the isolation of spatter within clusters 3 and 7. While not all the spatter was isolated by FDC in feature space, one could label spatter in the projection for better results. This could be done in the Kylberg experiments too. We wrap up by evaluating the effect of coupon positioning for this experiment by coloring the projection by coupon row, column, and index in Figure [4.67](#).



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.67: UMAP visualization of the Thermal Tomography ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

Similar to the Kylberg-WCNN experiment, the back row coupons 1, 2, and 3 tend to compose coarser regions in the projection with middle row coupons 4, 5, and 6 residing more in the fine region. In addition, coupons are well distributed in the smooth region. It is possible coupon-positioning is an influential factor we need to address in the analysis of thermal tomography. We use our final experiment to confirm some of these observations in

the next section.

ALOT-WCNN

Finishing with ALOT-WCNN, we cluster the thermal tomography feature space using a noise threshold of 0 which produces 15 clusters, the fewest out of all experiments. We relabel these clusters sorted by size in Figure 4.68.






Figure 4.68: Thermal tomography cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.


This distribution is not very uniform. Most of the data is grouped into clusters 13, 14, and 15. We expect these clusters to have a low textural homogeneity as their sizes imply they must capture multiple texture categories and/or transitions between them. We expect clusters 3 to 12 will be useful for isolating categories that are not too specific nor too broad. In Figure 4.69, we randomly sample and label images from each cluster.




(1) Solid; dark specks




(2) Solid; dark spots




(3) Smooth; edge/corner gradients




(4) Fine; gradients




(5) Solid; dark spots




(6) Solid bright border; dark relief



(7) Fine



(8) Fine; edge reliefs



(9) Smooth; bright border; grains

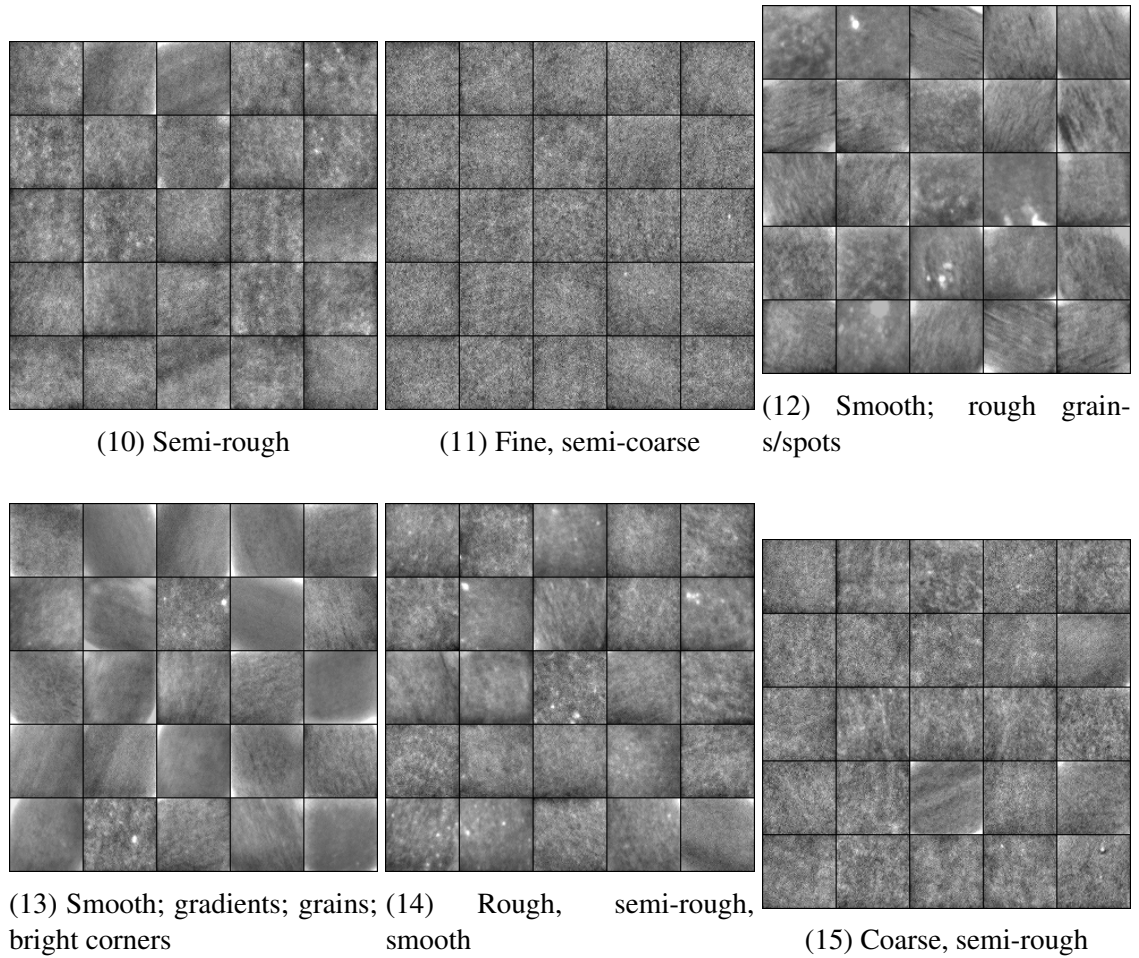


Figure 4.69: Normalized images sampled from each cluster in the thermal tomography ALOT-WCNN feature space.

This clustering still segments major groups, but there is more confusion. The distinction between coarse, semi-rough, and rough is not as clear regarding clusters 14 and 15. Unlike the ALOT-TCNN experiment, we do not observe any pure spatter clusters. We suspect most spatter is contained within clusters 12 and 14. Despite having fewer clusters, over-saturated samples still split into clusters 1, 2, and 5. Additionally, similar fine samples differing in reliefs split among clusters 4, 7, 8, and 11.

In general, we do not observe any new clusters or better qualities in this clustering. We proceed onto the analysis for the UMAP projection of this feature space colored by cluster in Figure 4.70.

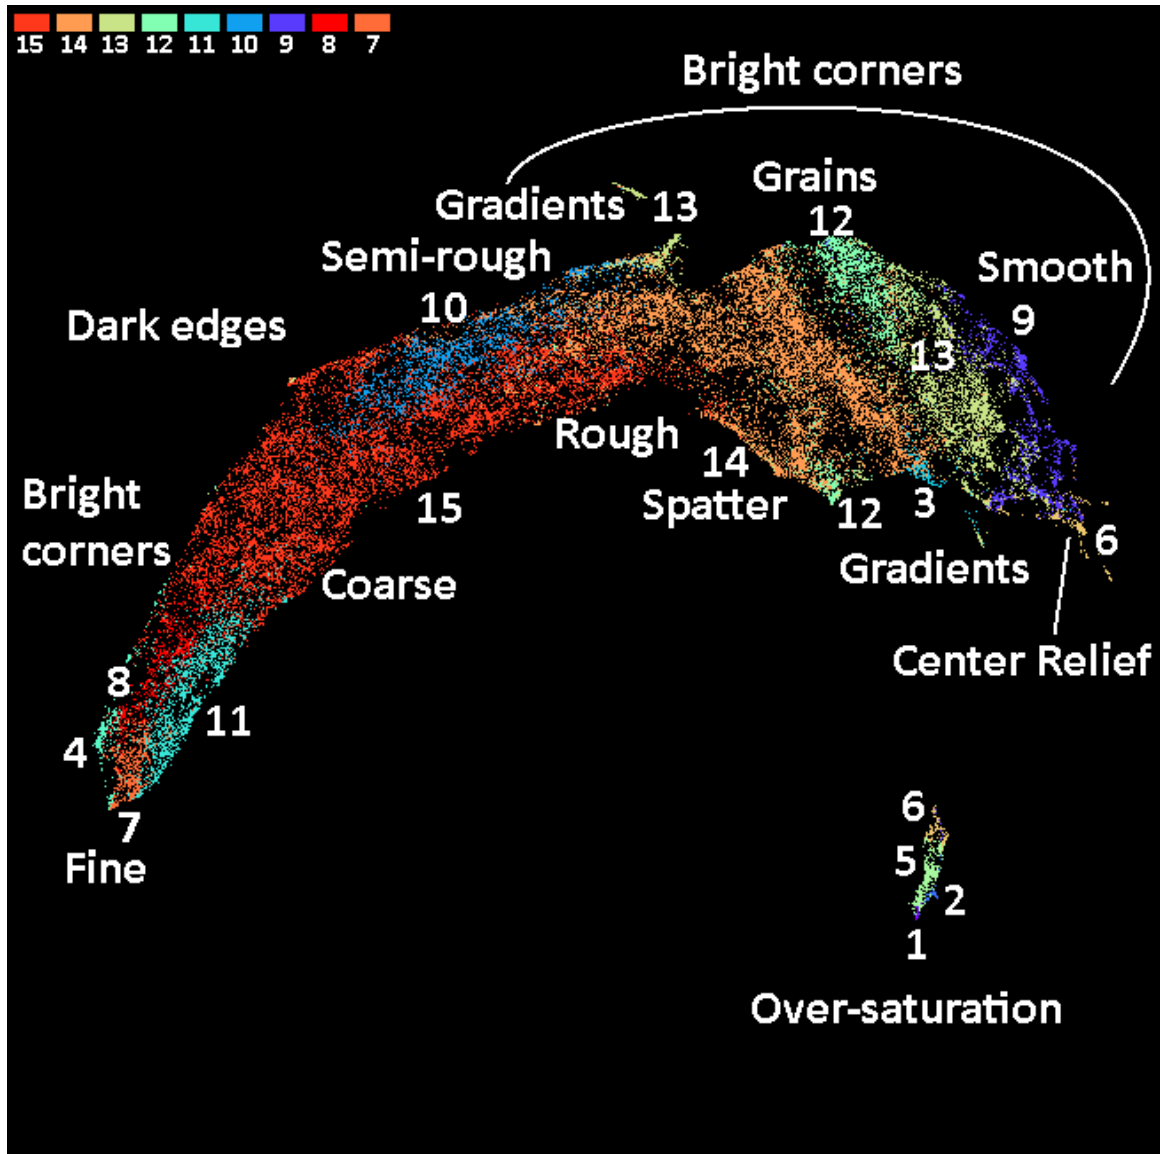
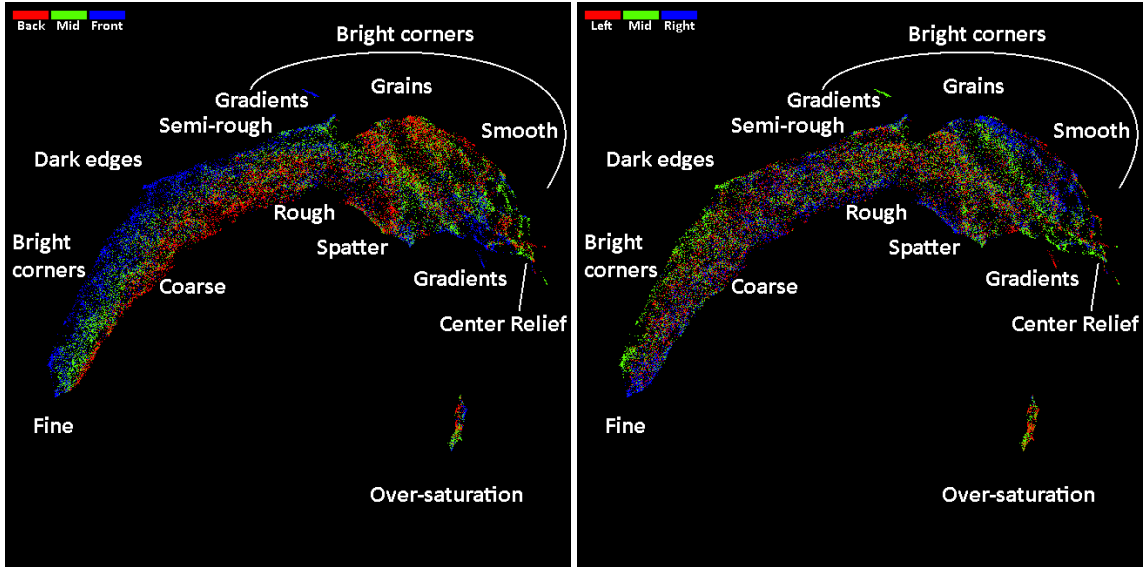


Figure 4.70: UMAP visualization of the Thermal Tomography ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.


The shape is quite different. Typically, we saw the fine, coarse, and rough regions forming a larger wider region in the projection, but this experiment shows them distributed along the projection tail. The topology is similar from a global scale, but along the transitions from fine to rough, the upper and lower halves segment by bright or dark features along the border. In this upper half of the tail, the fine region still wraps around to the semi-rough region. The spatter region is primarily contained in clusters 12 and 14 along

the bottom of the projection head up to cluster 3, but it does not cleanly break away. Some evidence of splitting is visible. We finish our analysis of this experiment by assessing the influence of coupon positioning by coloring the projection by coupon row, column, and index in Figure 4.71.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.71: UMAP visualization of the Thermal Tomography ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

This projection shows there is some influence of coupon positioning within non-

smooth texture categories. We observe more edge and gradient features along the top of the tail where front row coupons 7, 8, and 9 reside. This could be a result of lighting or cropping. If cropping captures too much of the coupon edge, a dark or bright gradient will be captured; however, when cropping excludes more of the border, it is less likely for these features to be included and possibly magnified from normalization. Once again, middle row coupons do not extend far into the rougher regions but reside more in the fine region. This phenomenon was observed in previous experiments include those on the visible recoat modalities.

This concludes our analysis of the ALOT-WCNN experiment. It primarily confirmed previous observations but also better exposed how coupon positioning was influencing the organization and segmentation of clusters and the projection.

Summary

In the thermal tomography data, we sufficiently clustered general categories of thermograms ranging from fine, coarse, rough, and smooth. Among these, we saw additional separation at transition areas within the projections and subcategories exposing over-saturation, grains, gradients, reliefs, and spatter. The spatial arrangement of clusters in the projections were consistent across all experiments.

Coupon positioning had a minor influence on the thermal tomography feature space structure similar to that in the recoat post-melt analysis. Coupons were grouped together row-wise within non-smooth categories. Middle row coupons represented finer thermograms more often than rougher ones which were populated primarily with back row coupons 1, 2, and 3. The cause could be related to lighting. Edges and corners were highlighted or recessed more in front row coupons 7, 8, and 9. We believe imprecise cropping contributed to the inclusion of these features.

From these results, we propose a few modifications. Over-saturation was more common for this modality, so it would be beneficial to tune the camera settings or extend the bit depth

of the thermal tomography images to prevent this from occurring. Tighter cropping to the bulk of each coupon would probably eliminate capturing border phenomena, but could help remove the influence of varying qualities of crops between coupons.

In these experiments smooth clusters contained high amplitude features such as bright or dark corners, edges, spots, and gradients. Larger clusters primarily captured fine, coarse, and rough thermograms which were best split into transitions and subcategories by the Kylberg-TCNN texture representation. Although the ALOT-TCNN clustering was not generally as homogeneous compared to the Kylberg clusterings, it was the only experiment to cleanly isolate spatter. This spatter segmentation excluded many rougher samples with abundant, smaller spatter. The ALOT-WCNN continues to show no advantage over other experiments in this modality and visible recoat. We propose further research focuses on the Kylberg-TCNN clustering because several categories, transitions, and special cases were isolated in clusters with high textural homogeneity. The ALOT-TCNN clustering could also be included for its segmentation of spatter.

In summary, the thermal tomography data was reasonably free of issues and segmented cleanly by general roughness categories. Some our goals were not fully achieved; they are the following. Cold spots were too scarce in the build to be isolated, but the segmentation of reliefs and gradients demonstrated a similar capability. Spatter segmentation was the most challenging with only some instances isolating in one experiment clustering. In the next section we analyze the LWIR data.

4.2.4 LWIR

The long-wave infrared (LWIR) data consists of peak-temperature heat maps from each coupon on each layer. The range of anomalies for this modality and preprocessing is still being explored, so a list cannot be reported at this time, but there are a few features that appear in the dataset. The most common feature is the presence of lines in nearly all samples. These lines reveal the laser scan path and are emphasized more after normalization. Spatter may still be manifested in heat maps as brighter blobs, but this is assumed to appear less often as the melt pool more commonly produces the hottest temperatures and the brightest pixels in each image. Texture washout via over-saturation should not occur because this modality uses a camera with a wide operating temperature range that extends beyond a practical maximum temperature.

In this section, we perform clustering to group similar heatmaps. This clustering would be beneficial if a few large groups with high textural homogeneity can be isolated. We begin by assessing how many clusters the LWIR feature space can break into while varying the noise threshold of FDC. We display the number of clusters per threshold for each experiment in Figure 4.72.

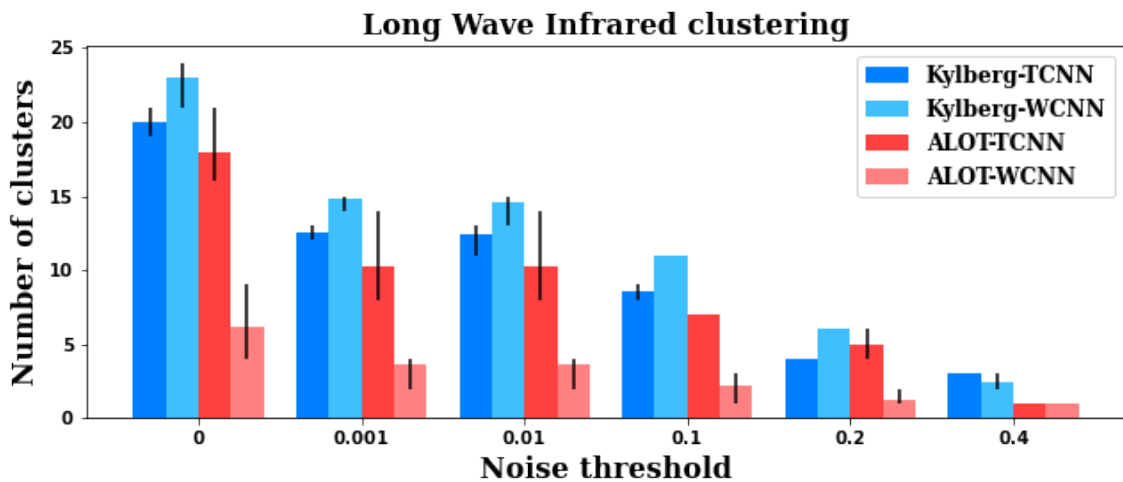


Figure 4.72: Plot of the average number of clusters the LWIR data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.

From previous modalities, the ALOT-WCNN experiments showed no advantages. As the ALOT-WCNN produces significantly fewer clusters than the other experiments, we suspect this experiment will produce the worst results. The number of clusters produced in the ALOT-TCNN experiments varies most indicating some clusters are not distinct and depend on initialization. The Kylberg experiments both produce more clusters than the ALOT experiments on average and are expected to yield a finer clustering. Finally, we observe at higher noise thresholds the number of clusters diminishes significantly implying these feature spaces are more uniformly dense than those from previous modalities. In the following experiment sections, we dive into a clustering and assess its contributions.

Kylberg-TCNN

Starting with Kylberg-TCNN, we select a noise threshold of 0 which produces 19 clusters. We relabel these clusters sorted by size and report their sizes in Figure 4.73.





Figure 4.73: LWIR cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.

From this distribution, it is clear that there are a few general groups captured by clusters 14 to 19. The smaller clusters 1 to 13 may be near-identical or rare instances. For visual assessment, we randomly sample images from each cluster and label them by textu-


ral description in Figure 4.74.



(1) Coarse lines; full dark border

(2) Smooth lines; large bright/dark corner/edge


(3) Smooth lines; full dark border



(4) Crisp lines; dark corner

(5) Smooth lines; bright/dark corner/edge


(6) Blurred lines; full dark border




(7) Smooth lines; bright corner

(8) Blurred lines; spotlight gradient


(9) Semi-coarse lines; dark corner




(10) Rough blurry lines; bright spots (11) Smooth lines; dark corner/edge (12) Smooth lines; bright corner/edge



(13) Smooth semi-coarse lines; dark corner (14) Smooth faded lines; bright corners (15) Smooth lines; bright spots/corner



(16) Rough blurry lines; bright/dark spots (17) Smooth semi-rough lines; bright corners (18) Smooth and coarse lines; bright/dark spots



(19) Crisp and coarse lines;
bright spots

Figure 4.74: Normalized images sampled from each cluster in the LWIR Kylberg-TCNN feature space.

As mentioned earlier, most of the laser scan lines are visible in each sample. Clusters vary depending on the roughness of these lines and border features. The smaller clusters 1 to 13 isolate specific instances of broader groups well. For example, clusters 1, 2, 3, and 6 isolate samples with dark borders. Cluster 8 captures blurry samples with spotlight gradients. These gradients might be induced by spatter. The center sample is one of a couple rare instances of data loss, thus it is left out of this discussion.

The major categories are spread across larger clusters 14 to 19. In clusters 14 and 15, heat maps with smoother lines and bright corners are isolated. Rough blurred lines with bright spots appear in cluster 16. The gap between smooth and rough blurry heat maps could be bridged by cluster 17. Samples from cluster 18 are more coarse with less blur. Finally, cluster 19 captures a mix of crisp and coarse heat maps. This cluster overlaps with cluster 18. In Figure 4.75 we color a UMAP projection of this feature space by cluster to visualize how clusters organize and overlap.

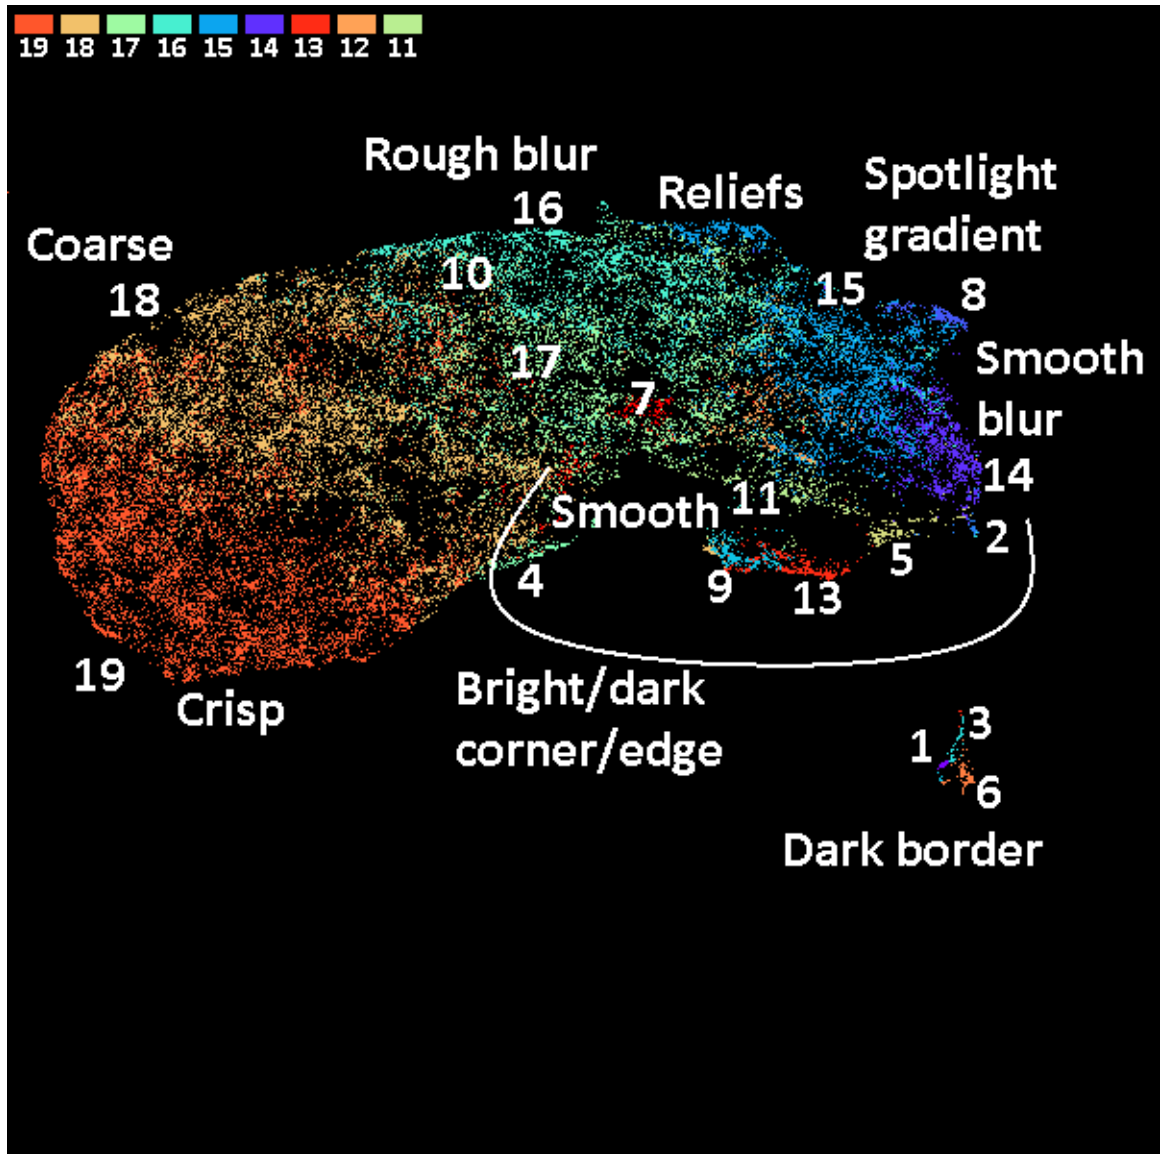



Figure 4.75: UMAP visualization of the LWIR Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

As seen in most modalities, the projection depicts a single cluster connected by textual overlap with similar semantically homogeneous regions. We observe a transition around the projection cluster passing through crisp, coarse, rough, smooth, and back to crisp. Samples get blurry when passing through the upper right region of the projection. The lower right region houses samples with significant bright/dark corners/edges. A small region breaks away from the other samples. These are heat maps with central dark borders. A


region at the top of the projection composed of cluster 15 is labeled with reliefs. These samples are dark and smooth within the center of the coupon, but their border is bright. These instances were not isolated into their own cluster but could in other experiments. We discuss the distribution of clusters. Notice cluster 19 extends through the coarse region that cluster 18 dominates. This was observed in the sampled images from cluster 19. Cluster 17 captures smoother heat maps residing below cluster 16 and extending towards clusters 11 and 15 and the transition between clusters 16 and 15. Clusters 9 and 13 pull away from the major projection cluster capturing semi-coarse lines with dark corners. Cluster 13 being smoother touches cluster 5 since it isolates dark corners and smooth lines. Extending off the blurry cluster 15 is cluster 8 whose samples are blurry with a significant gradient around highlighted spots. Smooth cluster 14 transitions into cluster 2 with a stronger emphasis of dark border features. Following in this direction leads to clusters 1, 3, and 6 which have completely broken away from the major cluster. These clusters isolate samples with full dark borders.

Overall, this clustering meets our goal, but the homogeneity of the larger clusters could be better. The distinction of significantly bright or dark spots is still a strong feature for segmentation. Unlike previous modalities, a blurry category had not been established until now. There are some concerns regarding coupon positioning. In the raw LWIR images before perspective calibration, there was a blur across the focal plane. This blur resulted in left column coupons 1, 4, and 7 being blurry while right column coupons 3, 6, and 9 were sharper. This modality has a lower resolution than the others, so samples had to be scaled up which could have introduced some blurring to all coupons. To determine whether coupon positioning has influenced this clustering, we color this projection by coupon row, column, and index in [Figure 4.76](#).



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.76: UMAP visualization of the LWIR Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

From these visualizations, we observe there is a strong correlation between the coupon column positioning and the texture of LWIR data. This relationship demonstrates coupons in the left column are blurry while coupons in the right column are crisp. This confirms our initial observation from inspecting pre-calibrated images. Additionally, upscaling images did not equalize the influence of blurring across all coupons. Coupon index and row do not

show a significant trend, but there appears to be some absence of back row coupons in the crisp region. In the proceeding experiments we point out improvements, consistences, and differences.

Kylberg-WCNN

Next, we cluster the Kylberg-WCNN feature space with a noise threshold of 0 yielding 23 clusters, the most of all experiments. We relabel these clusters sorted by size and report them in Figure 4.77.





Figure 4.77: LWIR cluster sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.


This distribution depicts a more gradual incline in cluster sizes implying groups have split more evenly. Similar to the Kylberg-TCNN experiment, the largest cluster still captures over a quarter of the entire dataset. The significant jump at cluster 23 implies it will have less textural homogeneity among its samples compared to other clusters. Larger clusters 13 to 22 may be isolating more common or broad categories of heat maps while smaller clusters 1 to 12 probably contain rare or specific textural heat map patterns. In Figure 4.78 we assign a textural description to a group of randomly sampled images from each cluster for visual inspection.



(1) Blurry; dark central relief

(2) Blurred lines; full dark border


(3) Smooth lines; dark corner



(4) Blurred lines; spotlight gradient

(5) Blurred lines; spotlight gradient


(6) Crisp lines




(7) Coarse lines, dark corner/edge

(8) Coarse lines; full dark border


(9) Coarse lines; small bright spots




(10) Crisp semi-coarse lines; bright spots




(11) Coarse lines; dark corners




(12) Smooth lines; bright corners




(13) Coarse lines




(14) Crisp coarse lines




(15) Smooth semi-coarse lines; edge gradient



(16) Smooth blurred lines; corner/edge gradient



(17) Coarse lines



(18) Crisp lines

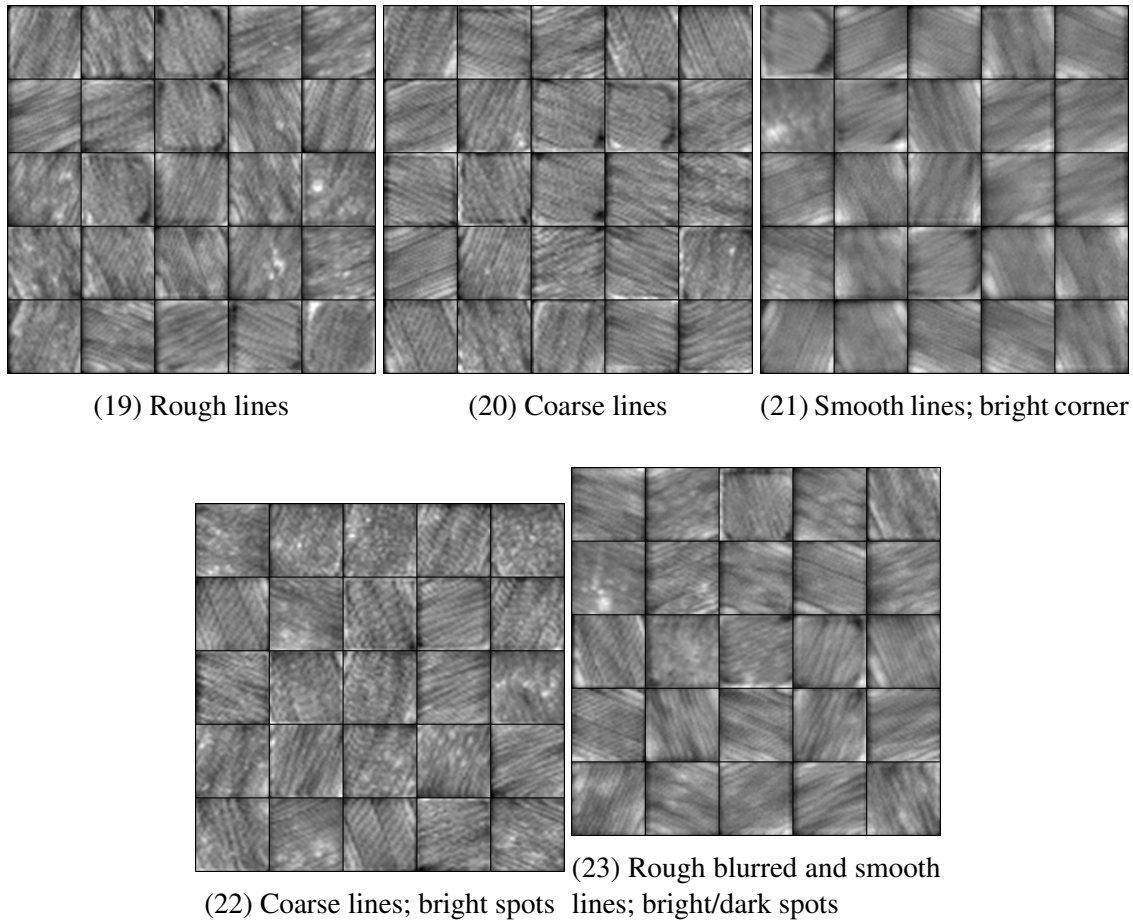


Figure 4.78: Normalized images sampled from each cluster in the LWIR Kylberg-WCNN feature space.

The images from these clusters depict reoccurring groups that were isolated in the Kylberg-TCNN experiment. We illustrate a few differences. There is a finer separation among crisp and coarse heat maps in clusters 6, 7, 9, 10, 13, 14, 17, 18, 20, and 22. The addition of two new categories are captured by clusters 1 and 19. Cluster 1 isolates instances with a dark central relief. These instances are rare, so their cluster happens to be the smallest. Cluster 19 segments rough heat maps with less blur. Unlike in the Kylberg-TCNN clustering, rough samples were grouped into cluster 16 regardless of blur. The largest cluster 23 mixes rough blurry samples with smooth ones instead of mixing crisp and coarse samples in cluster 19 from the Kylberg-TCNN experiment. These clusters are characterized in Table 4.5.

Texture	Clusters
Crisp	6, 10, 14, 18
Coarse	7, 9, 10, 13, 14, 15, 17, 20, 22
Rough	19, 23
Blurred	1, 2, 4, 5, 16, 23
Smooth	3, 12, 15, 16, 21
Border	2, 3, 8, 11
Gradients	4, 5, 16, 21
Center-relief	1
Spatter	17, 19, 22, 23

Table 4.5: Table of process phenomenon seen in Kylberg-WCNN LWIR clusters.

Levels of roughness and other features have been identified in each cluster. Spatter is more subtle in this modality and does not cover enough area per image to beat out the average feature activation; therefore it appears in rougher clusters but not as its own. This same phenomenon was observed in the thermal tomography. We proceed analysis by assessing whether the spatial arrangement of these clusters forms a similar category transition. Figure 4.79 depicts a UMAP projection of this feature space by cluster.

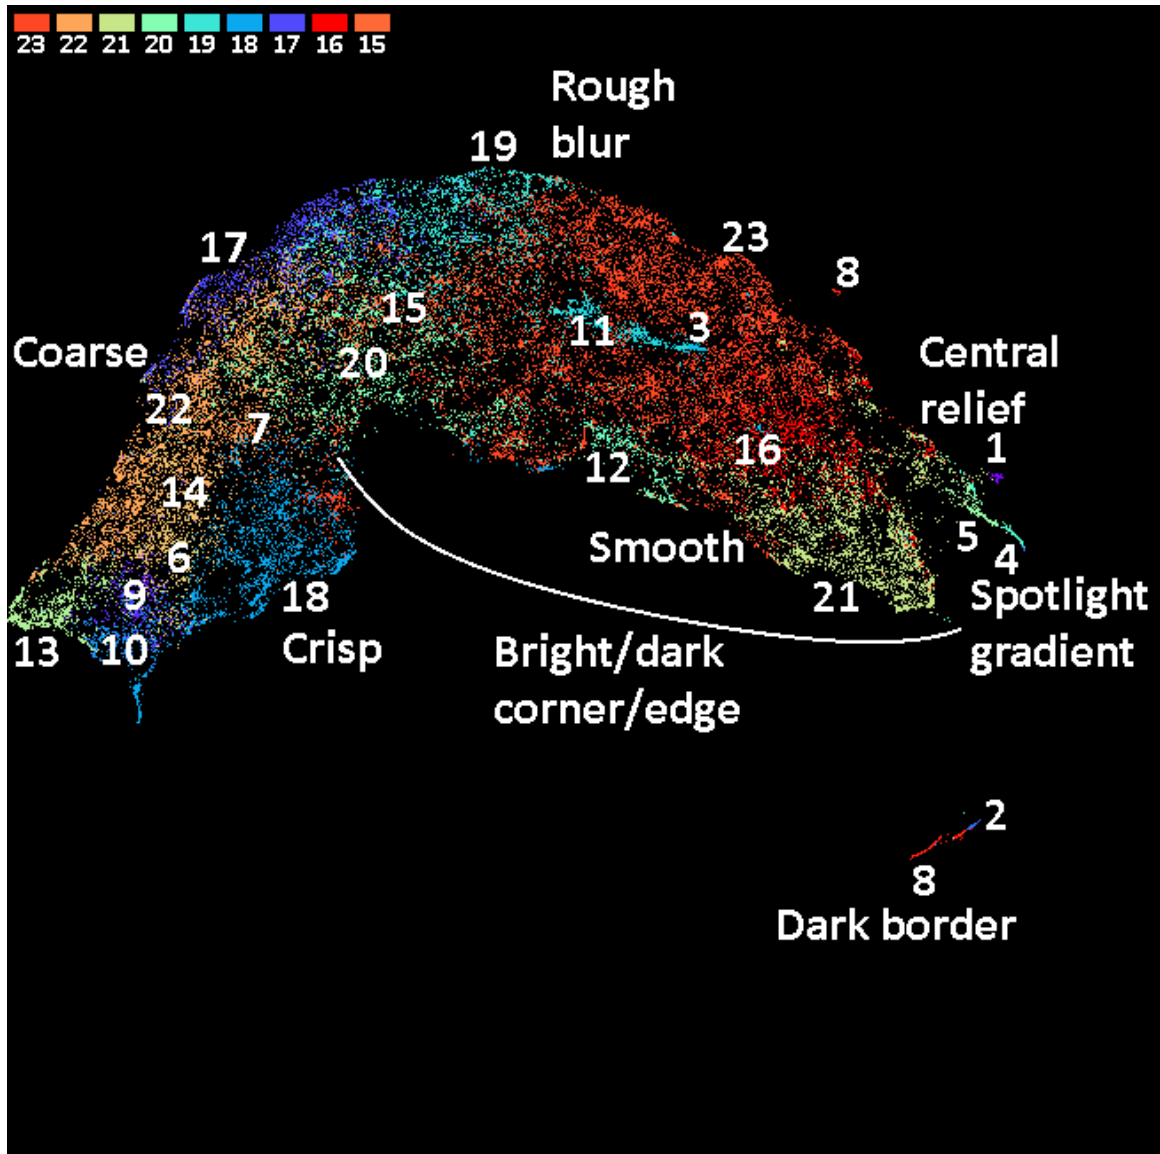
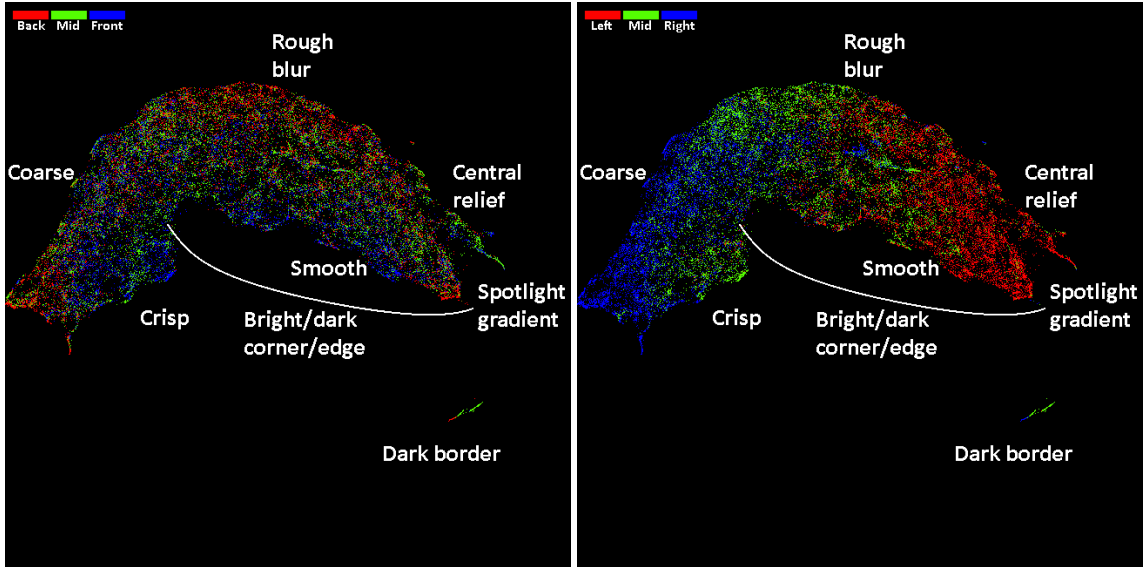


Figure 4.79: UMAP visualization of the LWIR Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.


We observe a similar projection forming a single cluster. The category transition from crisp through coarse, rough, and smooth is preserved. In addition, border features reside along the smooth-blurry region. Stronger examples of this are found in clusters 11 and 3 in the midst of cluster 23. It is not readily obvious why these two clusters are positioned on the projection in this way, but lines from these clusters are more coarse. In a similar manner to cluster 8 from the Kylberg-TCNN projection, clusters 4 and 5 extend off the

right side capturing blurry heat maps with spotlight gradients. Clusters 2 and 8 break away from the projection just how clusters 1, 3, and 6 did in the Kylberg-TCNN projection. We compare the effects of coupon positioning for this feature space by coloring the projection by coupon index, row, and column in Figure 4.80.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.80: UMAP visualization of the LWIR Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

From these visualizations, the column on the build plane still demonstrates a signifi-

cant relationship to the texture of the data. Looking at the row coloring, we observe a subtle bias where back row coupons are scarce in the crisp region. This was also observed in the Kylberg-TCNN visualization. The reason is not known but suspected to be an artifact in the data collection process.

This concludes our analysis of the Kylberg experiments. Both clusterings provided satisfactory results; however the effect of blurring along the column axis is a negative influence and might be corrected by adjusting the camera settings and positioning relative to the build plane. In the next section we dive into the ALOT experiments for comparative results.

ALOT-TCNN

We cluster the ALOT-TCNN LWIR feature space using a noise threshold of 0 to produce 15 clusters. We relabel these clusters sorted by sizes and plot their sizes in Figure 4.81.






Figure 4.81: LWIR cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.

This distribution is quite similar to the Kylberg-TCNN in that there is a clear distinction between large and small clusters. Clusters 12 to 16 could be isolating 5 broad categories while clusters 1 to 11 segment out specific or rare instances. To get a better understanding of the semantics behind these clusters, we label randomly sampled images


from each cluster in Figure 4.82.




(1) Blurred lines; spotlight gradient




(2) Dark central relief




(3) Smooth lines; bright corners; dark spot




(4) Coarse lines; dark border




(5) Smooth; full dark border




(6) Crisp semi-coarse lines; bright spots



(7) Coarse thick lines



(8) Smooth lines; bright blurry spots/corners



(9) Smooth thick lines; bright corner

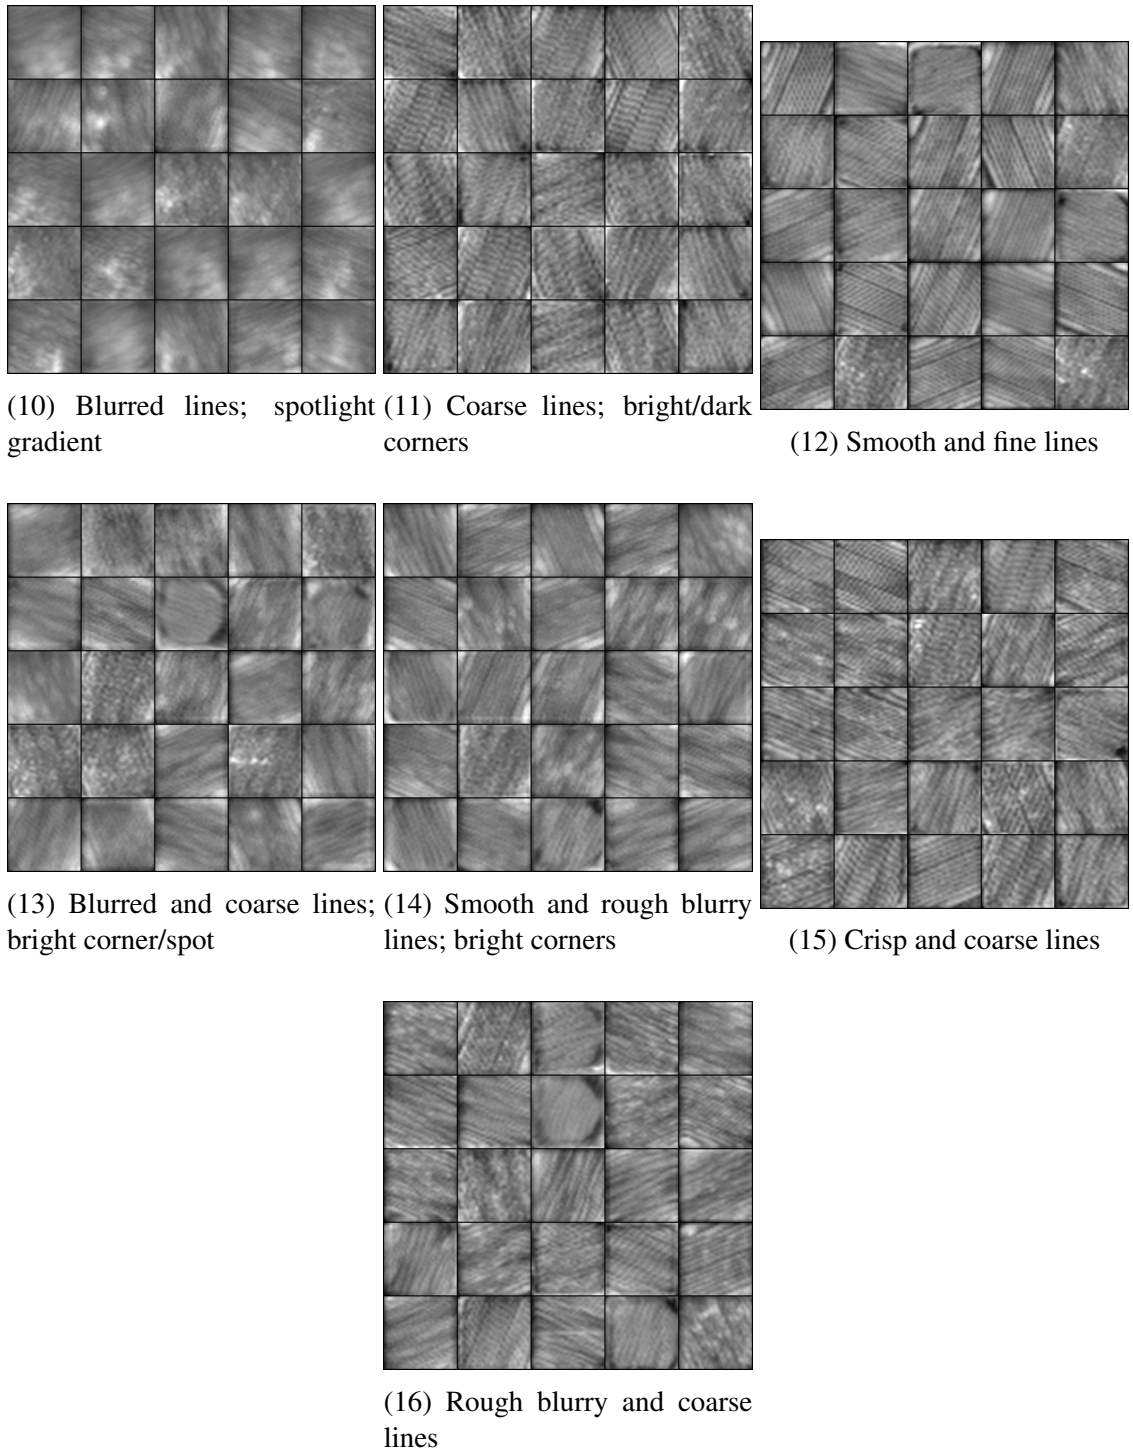


Figure 4.82: Normalized images sampled from each cluster in the LWIR ALOT-TCNN feature space.

This clustering is quite similar to those in the Kylberg experiments but offers two thick lines categories among smaller clusters 7 (coarser) and 9 (smoother). Many of the

smaller clusters isolate the same smaller groups which were seen previously, but dark border samples like those in cluster 8 from the Kylberg-WCNN are no longer being captured as purely. Some instances appear in cluster 4. In the larger clusters 12 to 16, general groups are intermingled. Cluster 12 contains a mix of crisp and smooth heat maps. The transition has likely been captured too. Cluster 13 groups blurry and coarse samples. Cluster 14 is similar, but mixes smooth and blurry rough categories which have been seen as neighbors from previous experiments. In a similar way, cluster 15 captures coarse and crisp samples. Finally, cluster 16 being the largest absorbs rough and coarse heat maps. This clustering implies the projection of this feature space follows a similar arrangement to those in the Kylberg experiments. We confirm this by coloring a UMAP projection of this feature space by cluster in Figure 4.83.

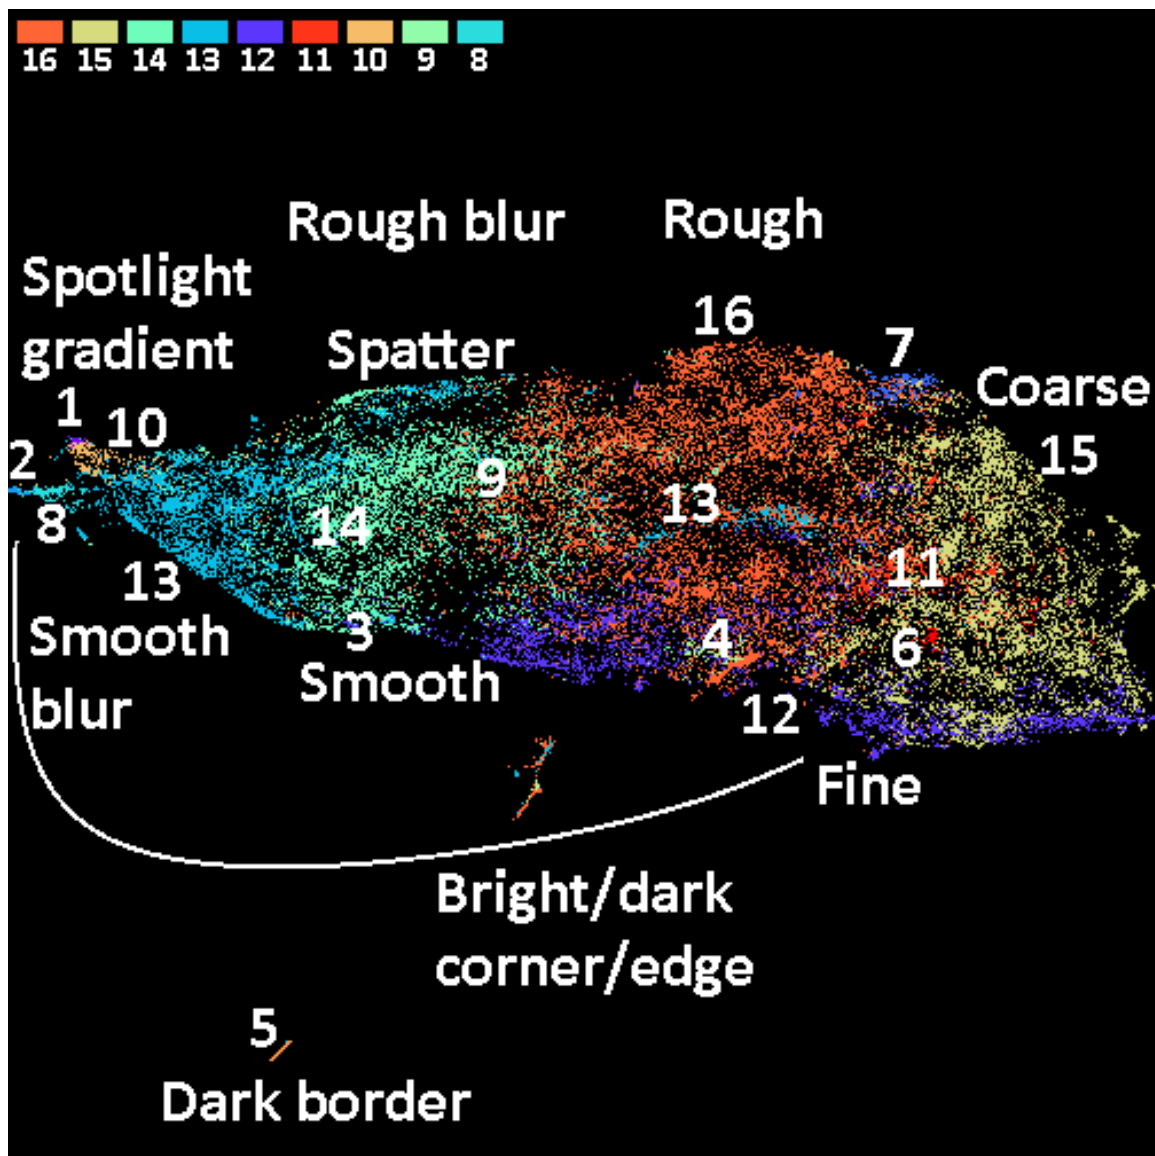


Figure 4.83: UMAP visualization of the LWIR ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

As expected, the projection depicts a single cluster whose texture category arrangement is nearly identical in the Kylberg experiments. Typically the full dark border samples break off into a single cluster in the projection, but in this projection those samples break off into two small strands. The larger of the two is not well isolated because it is primarily composed of cluster 16 followed by clusters 4 and 13. The smaller, further isolated strand is captured almost entirely by cluster 5. A new region has been labeled. The spatter region

is primarily located in the small horizontally-elongated cluster that splits off on the top of the projection tail. This region does not contain all instances of spatter for this data, but it is highly homogeneous. The spatter blobs are not as pronounced as in the thermal tomography data, but they are visible. To illustrate, we sample images from this spatter region and present them in Figure 4.84.






Figure 4.84: Normalized images sampled from the spatter region in the LWIR ALOT-TCNN feature space UMAP projection.

Overall this clustering has some small benefits but the distinction between broad categories captured among the larger clusters could be clearer. Additionally, the feature space projection isolates spatter into a region that could be labeled with high accuracy. We wrap up our discussion for this experiment by demonstrating the persistence of coupon positioning from projections colored by coupon row, column, and index in Figure 4.85.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.85: UMAP visualization of the LWIR ALOT-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

We observe the same trends hold. Left column coupon samples tend to be more blurry and smooth while right column coupon samples are crisp and coarse. Back row coupons are tend to populate the upper region half of the projection capturing rougher heat maps. This concludes our analysis for this experiment. In the next section we analyze the final experiment for this LWIR data.

ALOT-WCNN

Finishing with the ALOT-WCNN, we cluster the LWIR feature space using a noise threshold of 0 which yields only 4 clusters, the fewest of all experiments. We relabel these clusters sorted by size and plot their sizes in Figure 4.86.




Figure 4.86: LWIR cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.

This distribution shows a significant imbalance in cluster sizes that tend towards the trivial case of all points grouped into a single cluster. Clusters 1, 2, and 3 combined only make up a tenth of the dataset. We expect cluster 4 to be too broad for utility. In Figure 4.87 we label groups of randomly sampled images from each cluster for visual assessment.

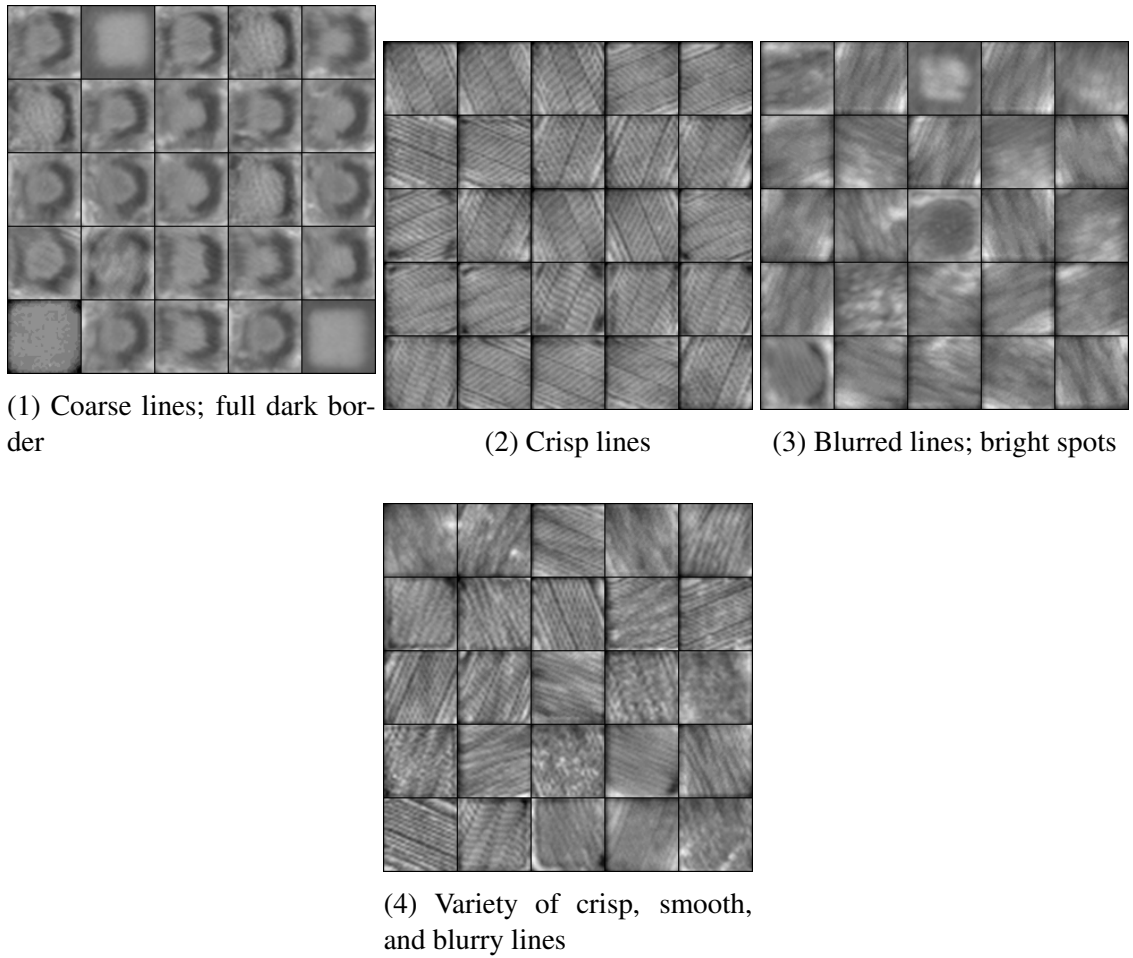


Figure 4.87: Normalized images sampled from each cluster in the LWIR ALOT-WCNN feature space.

As implied from the cluster sizes, cluster 4 contains a mix of textures that require further segmentation. The smaller clusters show some promise. Cluster 1 singles out a narrow group of heat maps with dark circular borders. This is the cluster that isolated well in the prior experiment projections. Cluster 2 captures a group of samples with crisp lines. Finally, cluster 3 spans a broad group of heat maps with blurry lines, spots, and reliefs. We show a UMAP projection of this feature space colored by cluster in Figure 4.88.





Figure 4.88: UMAP visualization of the LWIR ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.


In this projection we confirm the footprints of each cluster are proportional to their number of samples. The ends of cluster 4 splits their own clusters 2 and 3. Cluster 1 is a unique group and appropriately breaks away from the rest of the projection. Despite the poor clustering, the projection is comparable to previous ones in that the connectivity between texture categories is still preserved albeit the shape is a bit different. Similar to the ALOT-TCNN projection, the spatter region is cleanly segmented.

This experiment demonstrates a case where the performance of FDC degrades but the performance of UMAP does not. Our final assessment of this experiment looks at the influence of coupon positioning by coloring this projection by coupon row, column, and index in Figure 4.88.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.88: UMAP visualization of the LWIR ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

From these visualizations, we confirm the column-texture and subtle row-texture trends

hold. Overall this experiment does not provide a new benefit, clusters poorly, but confirms prior observations. In the last section we summarize our results for this modality.

Summary

In the LWIR data, we clustered categories of heat maps ranging from blurry to smooth, crisp, coarse, and rough. Among these, we saw additional separation at transition areas within the projections and subcategories exposing dark borders and spotlight gradients. UMAP visualizations of the feature spaces demonstrated spatial relationships of heat map categories were consistent across experiments. Furthermore, UMAP projections revealed the feature space of an experiment whose FDC clustering is poor may still preserve semantically homogeneous regions and transitions that could be manually labeled if the density of points is too uniform to segment clusters.

Coupon positioning had a strong influence on the LWIR feature space arrangement. We observed coupons were grouped together column-wise and row-wise but not individually. Unfortunately, this grouping occurred across heat map categories such that the distribution of textures were more blurry for coupons in the left column position and more rough for coupons in the back row position in the build plane. This makes it difficult to accurately identify the true phenomena manifested in the heat map without tailoring the approach to account for this non-uniformity in the build plane.

From these results, we recommend some future modifications to the data collection process and preprocessing. Blurring on left column coupons influenced the uniformity of texture representation across the build plane. We suspect this blur is induced by a misalignment between the build plane and camera focal plane. This leads to some coupons being out of focus in the collection of the LWIR data. To accommodate this we recommend finer tuning of the camera settings, focus, and positioning. The preprocessing strategy used the peak temperatures of layer frames to obtain an image per layer. This method produces brighter pixels along the scan path of the laser; therefore, lines are present in all images even if they

are subtle. We consider these lines to be an artifact that could interfere with clustering; therefore, we recommend developing a different preprocessing strategy that avoids embedding the laser scan path into the composite image.

Among these experiments, smaller clusters consistently captured specific or rare instances of heat maps including blurred spatter and dark central reliefs. Larger clusters primarily split heat maps by average roughness. These were best split by the Kylberg-WCNN texture representation which improved upon the Kylberg-TCNN. The ALOT-TCNN produced fewer clusters and more confusion among roughness categories within larger clusters. The ALOT-WCNN clustering split poorly, but its feature space was quite similar to the other experiments. We conclude the Kylberg-WCNN texture representation is best for clustering and further research.

In summary, the LWIR data had some obstacles, but we still achieved our goal of isolating general categories of heat maps. Unfortunately, these results are not applicable until the coupon positioning bias on texture is resolved. In the next section we analyze the spatter modality.

4.2.5 Spatter

The spatter data composes thousands of meltpool and ejecta emission images into a compiled image per coupon per layer. We developed this preprocessing strategy specifically for spatter to reveal the spatial distribution of meltpools and ejecta within a given area on the build plane. We have no prior knowledge of process phenomena that can be captured texturally using this preprocessing, but by design we expect brighter spots correspond to areas where the meltpool and ejecta were present more often and vice versa for dimmer spots. There are two obstacles we face from the nature of this data. First, image values need to be bounded, so we applied an activation function with a horizontal asymptote at the max pixel value. For coupons experiencing more ejecta and larger meltpools, texture washout can occur because an increase in pixel brightness associated with an increase in emissions decays as the pixel brightness tends towards to the max pixel value. This means that the distinction between large summed values is lost and mapped to the same pixel value which can produce areas in the final image with no variation. Second, pattern artifacts can be created when the periodicity of the frame rate and the laser hatching strategy overlap. Different processing parameters influence these artifacts with scan speed being the most influential.

In this section, we perform clustering to group similar spatter emission maps. The benefit of this clustering would be realized if we can segment at least two large texturally homogeneous groups. For example, one cluster may isolate smooth distributions of emissions while another captures rough distributions of emissions. We begin by assessing how many clusters the spatter feature space can break into while varying the noise threshold of FDC. We display the number of clusters per threshold for each experiment in Figure 4.89.




Figure 4.89: Plot of the average number of clusters the spatter data breaks into at different noise thresholds for each experiment. The range over 5 runs is indicated by a line on each bar.

From this plot, we observe the Kylberg-TCNN experiment produces over twice as many clusters as the other experiments. For most thresholds the Kylberg-WCNN, ALOT-TCNN, and ALOT-WCNN experiments produce nearly the same number of clusters. Whether these experiments produce the same clusters is discussed later. For the first time, the ALOT-WCNN tends to produce more clusters than the ALOT-TCNN and Kylberg-WCNN implying its clustering could be finer. In comparison to other modalities, the number of clusters each experiment produces is generally lower indicating there may be fewer categories of spatter emission maps in this build.

Kylberg-TCNN

Starting with Kylberg-TCNN, we select a noise threshold of 0 to produce 29 clusters, the most of the spatter experiments. We relabel these clusters sorted by size and plot their sizes in Figure 4.38.






Figure 4.90: Spatter cluster sizes of the Kylberg-TCNN feature space. Clustering noise threshold is 0.


This distribution splits groups of several clusters relatively uniform in size. These groups are illustrated in small clusters 1 to 16 averaging 250 samples, medium clusters 17 to 24 averaging 1500 samples, and large clusters 25 to 29 averaging 3000 samples. We expect the smaller clusters to isolate textures from a couple parameter sets, while medium and large clusters capture general groups. We inspect these clusters by labeling and displaying randomly sampled images from those clusters in Figure 4.91.




(1) Large lattice




(2) Smooth; reliefs




(3) Smooth; bright and dark edge




(4) Smooth; bright edge/spots




(5) Small lattice; bright edge




(6) Smooth; subtle spots; dark edge gradient




(7) Small lattice; bright edge; dark spots




(8) Semi-rough; edge gradient




(9) Smooth; bright corner; dark edge




(10) Subtle small lattice; bright edge/spot




(11) Large lattice




(12) Smooth; subtle spots; dark edge




(13) Subtle small lattice; large gradient; bright edge




(14) Smooth; bright edge; subtle spots




(15) Rough




(16) Smooth small lattices; central relief; bright edge




(17) Rough




(18) Smooth and solid; bright edge




(19) Smooth, subtly rough; dark edge




(20) Semi-rough




(21) Semi-rough




(22) Rough




(23) Smooth; gradients; subtle spots




(24) Smooth, semi-rough; spots




(25) Smooth, subtly rough; edge gradient



(26) Rough lattice



(27) Rough lattice



(28) Rough, semi-rough

(29) Rough, semi-rough

Figure 4.91: Normalized images sampled from each cluster in the spatter Kylberg-TCNN feature space.

This clustering depicts large clusters primarily segmenting by average roughness. Smaller clusters capture lattices and bright or dark edges, corners, and spots. These lattice artifacts vary in size depending on the scanning speed. Many clusters semantically overlap to the extent that some segmentation is unnecessary. For example, there is little to no semantic distinction between cluster pairs 20 and 21, 26 and 27, and 28 and 29. More examples exist, so we list clusters capturing different texture in Table 4.6 to demonstrate where overlap is occurring. Some phenomena like cold spots and severe spatter are scattered across too many clusters for concise listing and often appear in both smooth and rough samples.

Texture	Clusters
Smooth	2, 3, 4, 6, 9, 12, 14, 18, 19, 23, 24, 25
Semi-rough	8, 19, 20, 21, 24, 28, 29
Rough	15, 17, 22, 26, 27, 28, 29
Lattice	1, 5, 7, 11, 16, 22, 26, 27
Center-relief	4, 16

Table 4.6: Table of process phenomenon seen in Kylberg-TCNN Spatter clusters.

There are a few cases where better segmentation could be beneficial. Cluster 16 fails to separate small lattices from dark central relief samples likely because both include a bright edge and darker central area. Cluster 18 captures smooth samples, but some of

these are over-saturated and could be distinguished for that reason. Clusters 22, 26, and 27 capture rough emission maps. Some of their samples include lattices varying in levels of subtlety. This creates overlap in the embedded space making it difficult to draw the line between clusters. Like in the thermal tomography and LWIR, isolating spatter is still a challenge as many clusters contain some samples with at most a few small bright particles. This reoccurrence suggests a texture class featuring a few bright spots was absent from the training set. The experiment model will not learn features describing spatter if they are not present in the training set. Cold spots have a similar problem. Although the presence of cold spots is segmented between texturally similar clusters 5 and 7, the latter has noticeably sharper lattices which happens to include cold spots too. To understand how these clusters are arranged spatially, we color a UMAP projection of this feature space by cluster in [Figure 4.92](#).

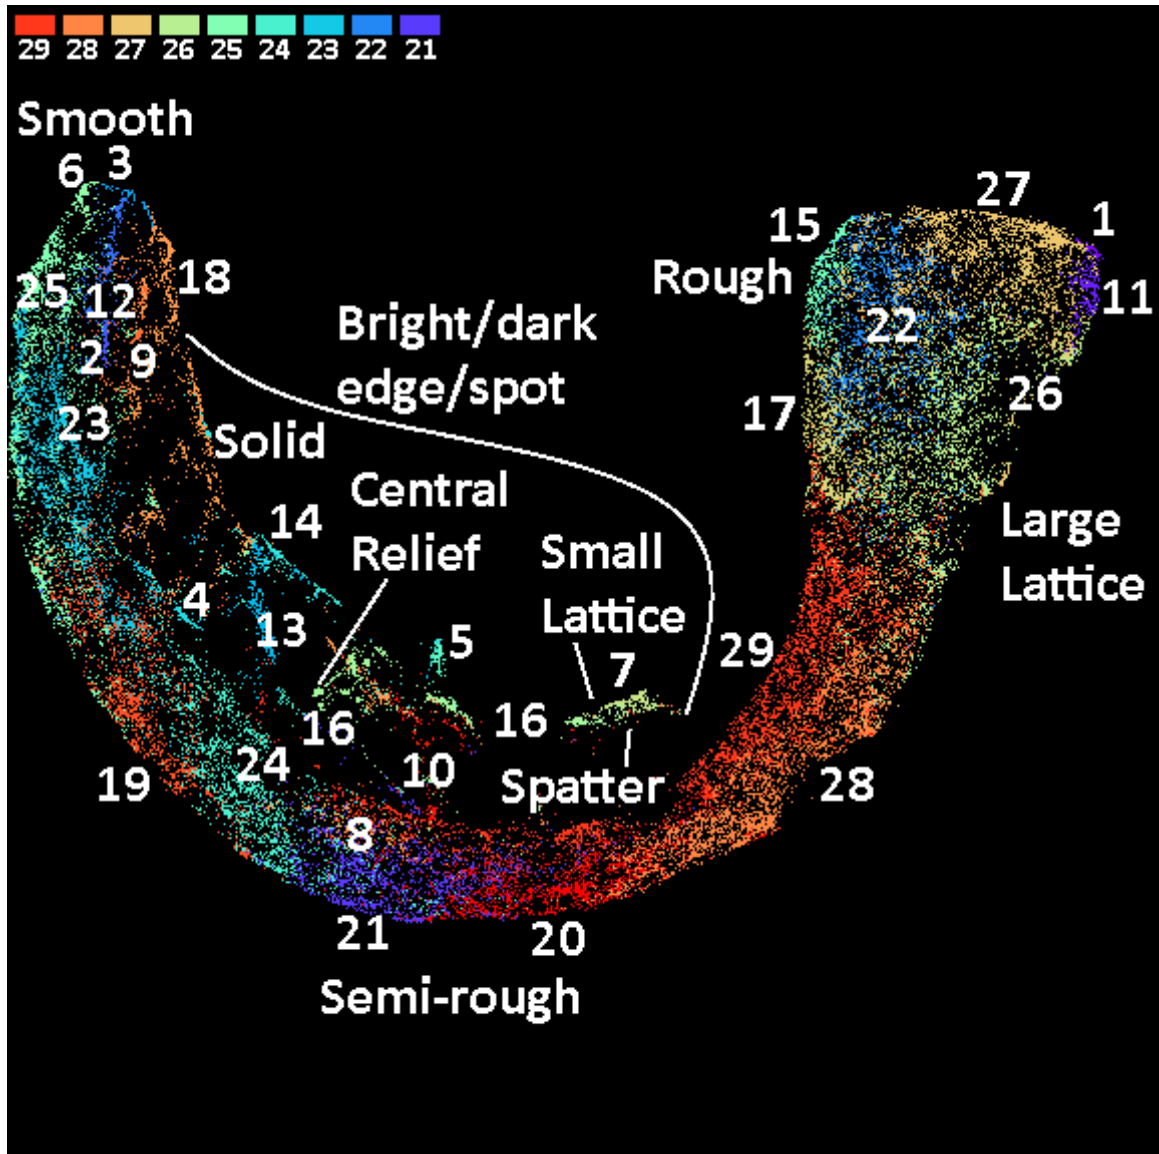


Figure 4.92: UMAP visualization of the Spatter Kylberg-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

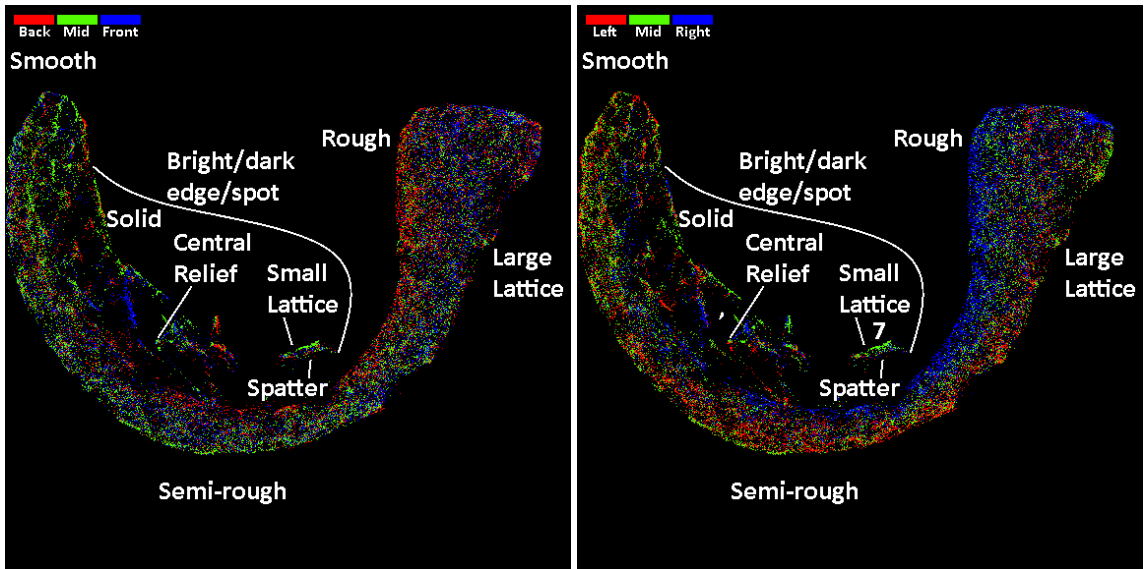
This projection depicts a single major cluster. It appears all of the LPBF sensor modalities from this build overlap to the extent that UMAP consistently projects LPBF feature spaces as a single large cluster with semantically isolated regions. These regions transition gradually from rough to smooth with many samples capturing large lattices in the rough region and many samples with bright or dark edges, spots, and corners in the smooth region. This general trend was observed in the other modalities, except coarse and fine regions do

not exist in this projection because there are no samples from those categories in this spatter data.

We make a few observations about the sparse regions in the projection. These regions are solid, small lattice, central relief, and spatter. Samples from these regions capture bright or dark edges, spots, and corners. The solid region is small capturing at most a couple hundred samples bright a brighter edge experiencing over-saturation due to abnormally high emissions. Many samples from the central relief, small lattice, and spatter regions are grouped into cluster 16. The small lattice region is primarily captured by cluster 7 and neighbors the spatter region. This spatter region only isolates a small subset of all the spatter samples.


Many clusters overlap in the projection while the distinction between some may not be immediately applicable. Consider clusters 1 and 11. They occupy the same region in the projection yet show no significant differences in their large lattice samples. The same argument can be extended to rough clusters 26 and 27 and their more subtle lattices. Although clusters 1 and 27 neighbor each other, it is useful that they are separated because the textural representation of a large lattice from cluster 1 should be distinct from a spatially in-homogeneous rough sample from cluster 27. In a similar manner, clusters 15 and 17 contain rough samples without lattices and could be combined but remain distinct from clusters 26 and 27 to preserve the distinction between rough samples with and without traces of lattices. Cluster 22 is a transition cluster containing rough samples with faint or missing lattices. Transitioning towards the semi-rough region, clusters 28 and 29, 19 and 24, and 8, 20, and 21 could be combined as they do not show significant differences. In the smooth region, clusters 23 and 25 could be joined. The other clusters are distinct enough to warrant their own label; however cluster 18 would be better split as it extends down to cluster 10 capturing less similar samples along the way.

Next, we visually assess whether texture is biased from coupon positioning by coloring the projection by coupon row, column, and index in Figure 4.93.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.93: UMAP visualization of the spatter Kylberg-TCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

From these figures depict subtle biases in coupon row and column. Middle row coupons tend to populate the smoother side of the projection. Right column coupons reside more along the concave side of the projection. In the rough region, this inner edge is composed of samples lacking large lattices. We look to other experiments first before making any additional conclusions. This concludes our analysis of the Kylberg-TCNN experiment.

Kylberg-WCNN

Next, we cluster the Kylberg-WCNN feature space with a noise threshold of 0 yielding 12 clusters. We relabel these clusters sorted by size and report them in Figure 4.94.






Figure 4.94: Spatter sizes of the Kylberg-WCNN feature space. Clustering noise threshold is 0.


This distribution exposes a gap between large and small clusters. Clusters 11 and 12 are large likely having lower textural homogeneity and corresponding to the smooth and rough categories. Clusters 1 to 6 may reflect some of the smaller clusters from the Kylberg-TCNN experiment. We expect clusters 7 to 10 capture broad texturally homogeneous groups. In Figure 4.95 we randomly sample images from each cluster and label them.




(1) Large lattice; dark spots




(2) Smooth; bright edge/strands edge; subtle spots




(3) Smooth; dark center; bright




(4) Smooth; bright edge; subtle spots




(5) Small lattice; bright edge




(6) Small lattice; bright edge; dark spot




(7) Smooth; bright edge



(8) Smooth; bright edge



(9) Smooth, subtly rough



(10) Smooth, subtly rough (11) Smooth, semi-rough (12) Rough, semi-rough

Figure 4.95: Normalized images sampled from each cluster in the spatter Kylberg-WCNN feature space.

In the smaller clusters, we observe many resembling those in the Kylberg-TCNN experiment. Cluster 1 captures samples from previous clusters 1 and 11. Clusters 3 and 4 appear as subgroups from previous cluster 16. Clusters 2, 5, and 6 correspond to previous clusters 4, 5, and 7, respectively. The medium-sized clusters are nearly identical to each other isolating smoother emission maps. Finally, cluster 11 captures smooth to semi-rough samples while cluster 12 captures semi-rough to rough samples. To illustrate the segmentation spatially, the clusters are colored on a UMAP projection of the feature space in [Figure 4.96](#).




Figure 4.96: UMAP visualization of the Spatter Kylberg-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.


The shape and topology of this projection is almost identical to the Kylberg-TCNN experiment. There are some minor differences. Within the rough region, large lattices reside along the concave edge of the projection instead of the convex edge. Additionally, some of these large lattices have broken off into their own cluster in the projection but are still assigned to cluster 12. Samples from this large lattice cluster came from clusters 11, 26, and 27 in the Kylberg-TCNN experiment. The dark central relief samples group tighter

and cleaner in the projection but are grouped into medium-sized clusters 9 and 10 instead of their own cluster. Spatter is scattered along the projection, so it does not have a labeled region.

We discuss the association of clusters to the regions. Cluster 12 captures the entire rough region mixing large lattices with spatially in-homogeneous rough emission maps. Some of these large lattices stem from this cluster into cluster 1. Cluster 11 spans the semi-rough region extending toward the subtly rough region captured by cluster 9 which blends into the smooth region in cluster 10. Cluster 7 and 8 are similar to clusters 9 and 11 but differ in the presence of bright edges causing them to spatially pull away from the rest of the projection. Clusters 5 and 6 capture small lattices. Finally, clusters 2, 3, and 4 occupy a region that is both smooth and semi-rough induced by the presence of bright or dark edges and spots.


Overall, this clustering loses distinction between some features such as the large lattices but captures general roughness categories with high homogeneity in as few clusters as possible.

We assess the influence of coupon position in Figure [4.97](#).



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.97: UMAP visualization of the spatter Kylberg-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

These visualizations cross validate three previous observations. There are no large regions in the projection that belong solely to a single coupon. The middle row coupons slightly bias populating smoother regions. In rough regions the right column coupons have a tendency to populate the side that is opposite to the large lattice domain; however some of the clearest examples of large lattices found near cluster 1 are from these right column

coupons. The impact and logic behind the positioning-texture bias is still unclear, so we proceed analysis with the ALOT experiments.

ALOT-TCNN

We cluster the ALOT-TCNN spatter feature space using a noise threshold of 0 which produces 12 clusters, the same as in the Kylberg-WCNN experiment. We relabel these clusters sorted by sizes and plot them in Figure 4.98.



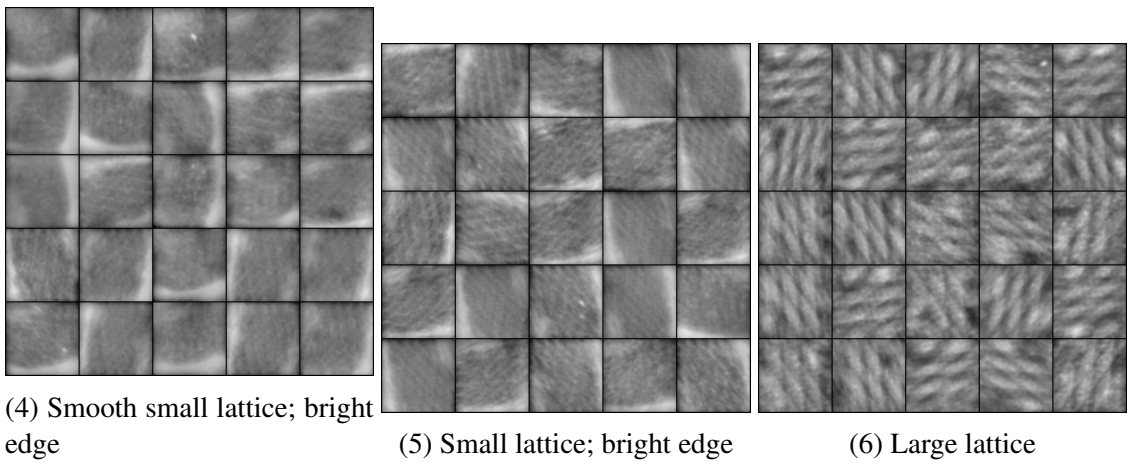



Figure 4.98: Spatter cluster sizes of the ALOT-TCNN feature space. Clustering noise threshold is 0.

This distribution tends towards the trivial case of all data being grouped into one cluster. It appears quite different from the Kylberg-WCNN cluster size distribution implying the semantics of the clusters are different despite the number of clusters being the same. In Figure 4.99 we show labeled images randomly samples from each cluster for visual assessment.






Figure 4.99: Normalized images sampled from each cluster in the spatter ALOT-TCNN feature space.

The small clusters are similar to those in the Kylberg experiments. Samples from clusters 1, 4, 5, and 10 correspond to those in clusters 1, 5, 6, and 3 from the Kylberg-WCNN clustering, respectively. Clusters 2, 3, and 7 draw samples from clusters 14, 18, and 13 in the Kylberg-TCNN experiment, respectively. Some of the larger clusters are similar too. Clusters 8 and 9 contain samples from smooth clusters 8, 9, and 10 in the Kylberg-WCNN clustering. Although cluster 11 is significantly smaller than Cluster 12, its image sampling implies lower homogeneity within the smooth bright edge region. As cluster 12 captures most of the data, both smooth and rough samples are grouped together. This segmentation behavior of ALOT-TCNN feature spaces has been confirmed in recoat post-melt, thermal tomography, and LWIR clustering analysis; however this texture representation has also shown to group instances of spatter better in the projection. To verify if this is the case, we color a UMAP projection of this feature space by cluster in Figure 4.100.




Figure 4.100: UMAP visualization of the Spatter ALOT-TCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.

This projection depicts a different shape but similar topology. The transition of regions from large lattices to rough, semi-rough, and smooth is still preserved. Additionally, the sparse region containing bright/dark corners, spots, and edges resides alongside the smooth and semi-rough regions. Some of the less common regions are better isolated in this projection. Small rough lattices sit between the transition from large lattices to rough emission maps. These are different from the small lattices captured by clusters 4 and 5

because those feature a bright edge. Spatter groups better stemming off slightly from the rough region. The samples are abundant in these region and highly homogeneous in semantics. We demonstrate this by showing an image sampling from this region in Figure 4.101. The solid region cleanly splits away in the projection, but its samples are still grouped with non-saturated samples in the cluster 3 region neighboring cluster 11. Finally, nearly all central relief samples are grouped well in the both the projection and in cluster 10.


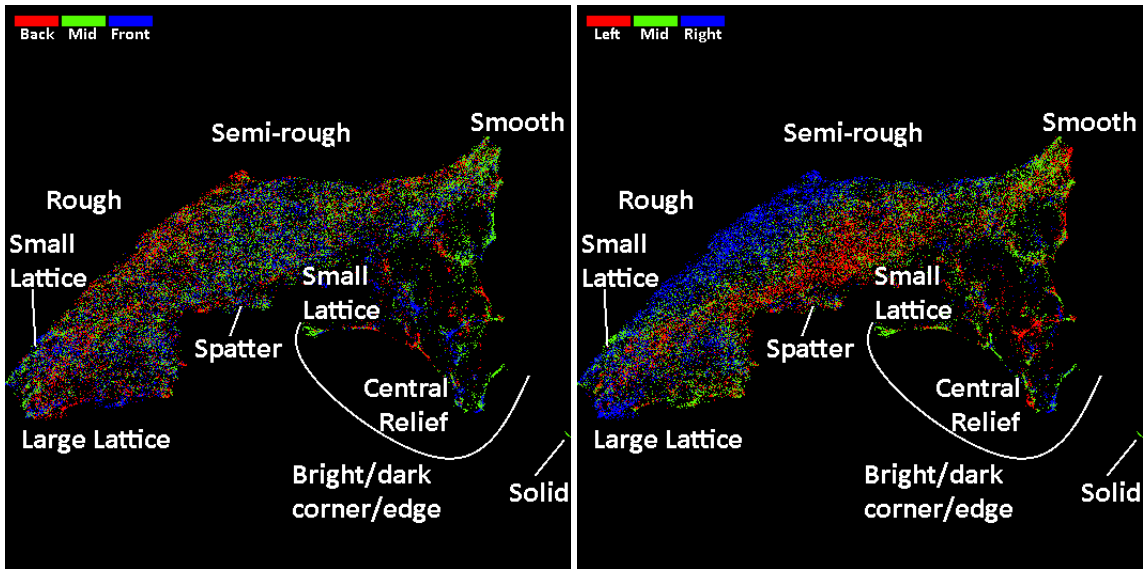



Figure 4.101: Normalized images sampled from the spatter region in the Spatter ALOT-TCNN feature space UMAP projection.

Overall, this projection shows improvements over the Kylberg projections; however the clustering only has a couple minor benefits such that cluster 10 that do not out-weight the low homogeneity of cluster 12. In Figure 4.102, the projection is colored by coupon row, column, and index to provide more evidence for the discussion of coupon positioning biases.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.102: UMAP visualization of the spatter ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

Similar to the Kylberg experiments, coupon index shows no bias on texture, middle row coupons tend to be a little smoother, and the right coupon column shows a clear separation from the other coupons along the rough and semi-rough regions. After comparing hundreds of samples from the right column coupons to the other coupons, we observed the right column coupons were slightly sharper at a low-level indicating images from the

middle and left column coupons may be subtly blurred. This concludes our analysis on the ALOT-TCNN experiment. In the next section we discuss the last experiment for the spatter data.

ALOT-WCNN

Finishing with ALOT-WCNN, we cluster the spatter feature space using a noise threshold of 0 yielding 13 clusters, one more than in the Kylberg-WCNN and ALOT-TCNN experiments. We relabel these clusters sorted by sizes and plot them in Figure 4.98.






Figure 4.103: Spatter cluster sizes of the ALOT-WCNN feature space. Clustering noise threshold is 0.


This distribution shows an improvement in uniformity among cluster sizes compared to those in the Kylberg-WCNN and ALOT-TCNN experiments. While many clusters are still small, we expect better textural homogeneity among the larger clusters 11, 12, and 13 given that this broader clustering has fewer clusters than the Kylberg-TCNN clustering. We visually assess these clusters by labeling randomly sampled images from each cluster in Figure 4.104.



(1) Small lattice; bright edge; (2) Smooth; bright edge; dark (3) Smooth; bright edge; dark
small spots central area central area



(4) Smooth; dark edge (5) Smooth and solid; bright (6) Smooth lattice; large gradi-
edge ent; bright edge



(7) Smooth; dark corner/edge (8) Rough (9) Rough, semi-rough

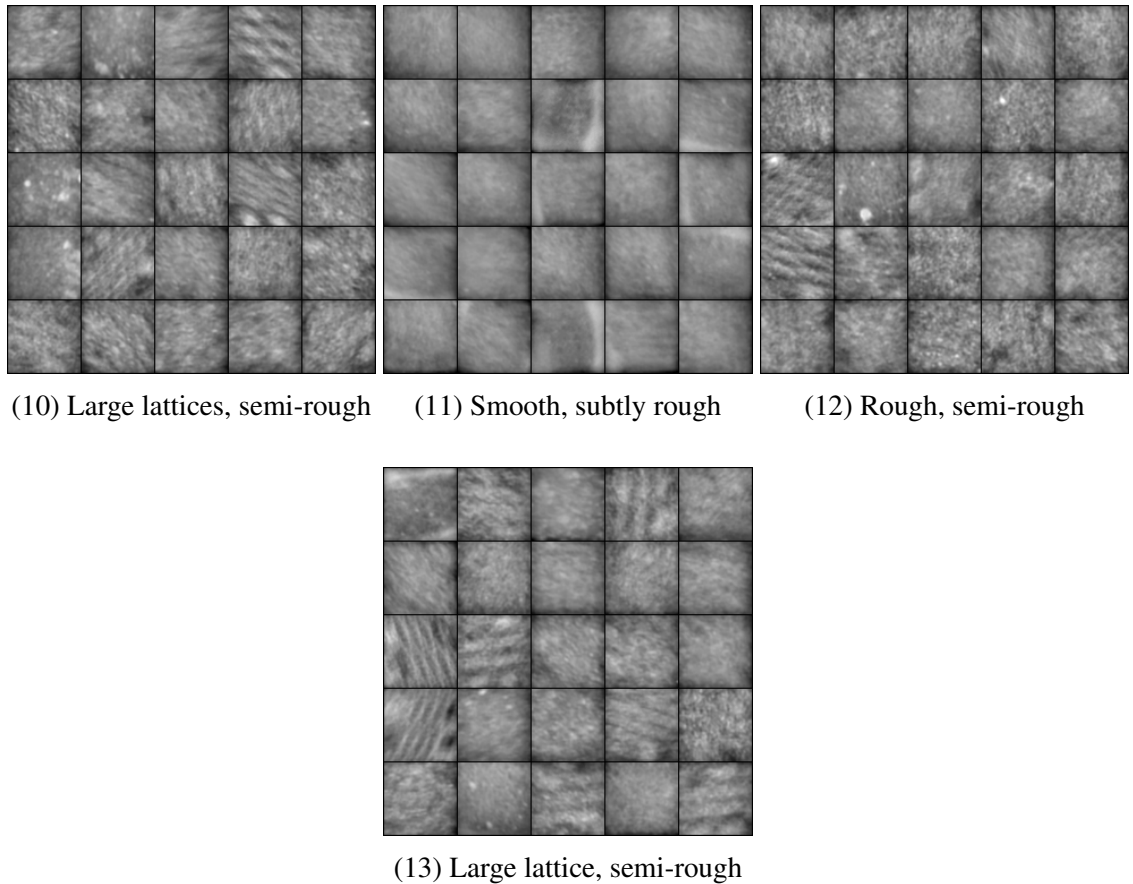


Figure 4.104: Normalized images sampled from each cluster in the spatter ALOT-WCNN feature space.

Despite the distribution looking better, the textural homogeneity of larger clusters 10, 12, and 11 appears lower. Additionally, there is no cluster that solely isolates large lattices nor small lattices with a bright edge and dark spots. Examples of these from the Kylberg-WCNN clustering are clusters 1 and 6. Some of these clusters have been seen before. Clusters 2 and 3 combined, 5, and 6 correspond to clusters 2, 3, and 7 from the ALOT-TCNN experiment. Clusters 8 and 9 correspond to clusters 15 and 17 in the Kylberg-TCNN experiment. Clusters 4 and 7 are both similar and capture samples like those in clusters 7 and 8 in the Kylberg-WCNN experiment. In Figure 4.105, we color a UMAP projection of this feature space by cluster to assess how these clusters are spatially organized.

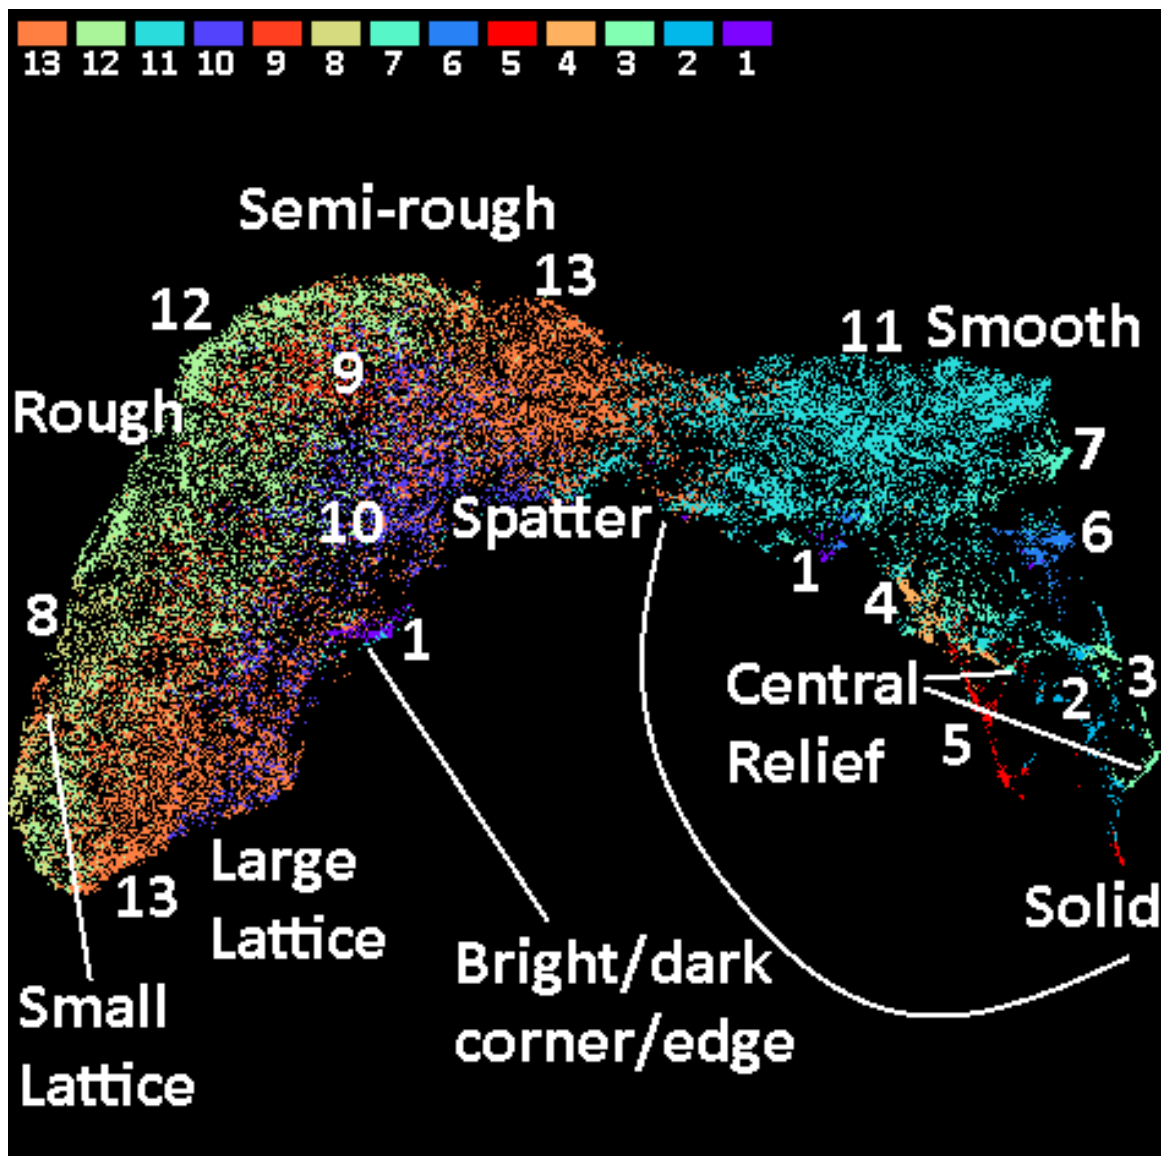



Figure 4.105: UMAP visualization of the Spatter ALOT-WCNN feature space colored by cluster. Minimum distance of 0.00 and 50 neighbors were used.


Once again, the shape is different, but the arrangement of texture categories remains the same. The large lattice region transitions through the small lattice, rough, semi-rough, smooth, and solid regions in that order. Like in the ALOT-TCNN projection, a region primarily containing spatter exists, but this region is not as pure. While the solid region segments better in the projection, it is split among clusters 2, 3, and 5. The central relief region does not group as well in both the projection and clustering.

We discuss how the clustering is spread along the projection. Clusters 13 and 1 are not contiguous in the projection. Large lattices and semi-rough samples are spatially separated in the projection but are grouped into cluster 13. The splitting of cluster 1 into the rough and smooth regions is reasonable. The samples from cluster 1 both have small lattices with a bright edge, but some are rougher while others are smoother. There is a significant amount of overlap among clusters 8, 9, 10, 12, and 13 within rougher regions. Cluster 11 dominates the smooth region lacking bright or dark corners, edges, and spots. Cluster 7 closely neighbors 11 while capturing dark edges/corners. Among clusters 2 to 5, various dark central reliefs are captured. Cluster 6 captures some samples with subtle small lattices, but those with dark spots have been grouped into cluster 13 beneath cluster 8 in the projection. Overall this experiment does not show a significant advantage. We conclude our analysis by verifying the bias of coupon positioning on texture holds. In [Figure 4.106](#) the projection is colored by coupon row, column, and index.



(1) Coupon row

(2) Coupon column



(3) Coupon index

Figure 4.106: UMAP visualization of the spatter ALOT-WCNN feature space colored by coupon features. Minimum distance of 0.00 and 50 neighbors were used.

These visualizations imply the coupon row and index do not have a significant influence on texture, but coupon column does. In the rough region, right column coupon samples have rougher low-level features while the left column samples have smoother low-level features. Additionally, middle column coupons appear to favor the smoother region. The cause behind these trends remains unknown. This concludes our analysis of the spatter

data. We summarize our findings for this data in the next section.

Summary

In the spatter data, we clustered categories of emission maps ranging from lattices to rough, semi-rough, smooth, and solid. Among these categories, transition regions within the projections exist, but transition clusters were scarce in most experiments. Without transition clusters, the utility of the clustering is degraded because the precision between roughness categories is lower. Subcategories including lattices and dark central reliefs were isolated in some clusters with bright spatter only isolating well in projections. UMAP visualizations of the feature spaces demonstrated roughness categories gradually transitioned into each other and sparse regions residing near the smooth region contains samples with bright or dark edges, corners, and spots.

Coupon positioning had a subtle influence on the spatter feature space arrangement. We observed coupons grouped together column-wise and row-wise but not individually. The column-wise bias arose from intra-category variations. This was primarily observed by the right column coupons having finer low-level features in rough samples. This influence is significant enough to beckon modifications. The row-wise bias spanned across texture categories such that middle row coupons were slightly smoother; however, this influence was too subtle to draw conclusions.

From these results, we recommend only a couple modifications to the data collection process and preprocessing. Rough right column coupon samples were finer at a pixel level. This could be an indirect result of a subtle blur on the other coupons induced by a small mismatch between the camera focal plane and the build plane. We suggest finer tuning of the camera settings, focus, and positioning to correct the bias. The preprocessing strategy applied an activation function onto the sum of layer frames to obtain an image per layer. This method produced dimmer pixels at points where it was less common for the camera to capture a frame when the melt pool was centered at those points; therefore, lattice patterns

are present in many images with some being more subtle. We consider these lattices to be an artifact that interfered with clustering; therefore, we recommend developing a different preprocessing strategy that avoids embedding these into the composite image.

Among these experiments, smaller clusters consistently captured specific or rare instances of emissions maps including spatter, lattices, and dark central reliefs. Larger clusters primarily split by average roughness. These were best split by the Kylberg-TCNN texture representation. The other experiments produced much fewer clusters. Both ALOT experiments did not segment roughness categories well, but their smaller clusters and projections isolated samples with higher homogeneity. The Kylberg-WCNN clustering performed well at isolating the general roughness categories in only a few clusters. We conclude the Kylberg-TCNN texture representation is best for clustering and further research.

In summary, the spatter data was clean and separated well by levels of roughness; therefore, we achieved our goal of isolating smoother and rougher spatter emission maps with an added bonus of other subcategories. This clustering can be labeled and used to describe the spatter emission map texture of new similar samples. In the next section we cluster coupon data from each layer using all modalities simultaneously.

Conclusion

In this thesis we developed an unsupervised approach to segment LPBF process monitoring imagery by texture using general texture representations. In this chapter, we summarize the data collection and preprocessing, the texture-driven clustering strategy, the contributions from the results, and our final thoughts.

5.1 Data Collection and Preprocessing

5.1.1 Texture Datasets

The texture datasets, Kylberg [16] and ALOT [6], were publicly available online. Both supplied features for learning texture embeddings. Kylberg was selected as baseline dataset that could concisely capture a wide variety of texture features in only 28 classes. ALOT was selected as a brute force approach with 250 classes to supersede the diversity of Kylberg in the event that some vital features in the LPBF data were missing in Kylberg.

5.1.2 LPBF Datasets

The LPBF data was collected from an experiment designed to produce a wide variety of process phenomena that could be captured via in-process monitoring without causing build failure. The build consisted of printing nine coupons in a 3-by-3 grid 3267 layers tall using

a total of 243 unique processing parameter combinations varying in laser power, scanning speed, hatch spacing, powder layer thickness, and laser spot size. Each coupon was sectioned into 27 equally-sized height segments each using a randomly assigned parameter combination to prevent any textural bias in the build plane.

Data was collected from four sensors, namely visible recoat (post-spread and post-melt), thermal tomography, long-wave infrared (LWIR), and spatter. Each sensor modality had its own unique preprocessing steps followed by camera calibration and standardization. These preprocessing steps are the following.

Visible Recoat

The visible recoat data was collected by taking two images per layer, one after spreading powder and another after melting it; therefore, no compiling step was necessary. Due to prior sensor damage, each image had a thick horizontal line of dead pixels obstructing coupons 4 and 5. This damage was mitigated by filling rows of dead pixels with the nearest neighboring rows, but these pixels were subtly darker. For this reason, we called them dead pixel residue.

Thermal Tomography

The thermal tomography data was collected in multiple frames while melting to describe the thermal topography per layer. These frames were summed together pixel-wise to create a composite image. The compiling process produced bright pixel artifacts in the composite image as a result of compounded small sensor noise. These artifacts were removed by replacing their pixels with the median of their unaffected neighbors.

LWIR

The LWIR data was collected in multiple frames while melting to model the cooling rates per layer. Composite images were constructed pixel-by-pixel from the peak temperatures across all frames. The compiling process produced dark line artifacts in the composite images where the hatch space sat between laser passes.

Spatter

The spatter data was collected in multiple frames while melting to record quantitative statistics of spatter particles emitted per layer. These were summed together pixel-wise followed by an exponential scaling function to create a composite image. The compiling process produced lattice artifacts in the composite images when the processing parameters and hatching strategy shared a periodic overlap with the frame rate of the camera.

5.2 Texture-Driven Strategy

The texture-driven image clustering strategy was split into a supervised part followed by an unsupervised part.

In the supervised part we adopted and modified two published CNNs (TextureCNN [4], WaveletCNN [27]) tailored for texture classification and trained each one on two public texture datasets, namely Kylberg and ALOT to obtain four texture embeddings for comparison.

In the unsupervised part after training all four models, the encoders were used to embed texture datasets and LPBF imagery from each sensor modality into compact feature spaces. The next step was to cluster these feature spaces using an appropriate clustering algorithm. We suspected the distribution of anomalies to nominal phenomena in the LPBF data would be imbalanced, so we used a density-based clustering algorithm (FDC) for its ability to

quickly find a non-user specified number of clusters of varying in size, shape, and variance. After clustering, we randomly sampled images from each cluster for visual assessment. To supplement analysis, we used UMAP for feature space visualization because of its speed and ability to better preserve global information. Feature spaces were visualized on scatter plots colored by cluster and labeled to demonstrate spatial relationships and transitions between clusters as well as mark regions belonging to general texture categories. Finally, the applicability of the LPBF data was empirically observed by finding biases in the texture distribution associated with the positioning in the build plane. This was achieved by coloring points in the projection that were associated with image samples taken from a specific coupon row, column, and index.

5.3 Contributions

The contributions from the results in this thesis are discussed separately per dataset starting with preliminary observations drawn from the analysis of Kylberg and ALOT followed by the conclusions from the analysis of visible recoat post-spread and post-melt, thermal tomography, LWIR, and spatter. Afterwards we summarize the comparison between experiments.

5.3.1 Kylberg

The analysis of Kylberg showed most classes could be isolated in both the clustering and projection. Spatially in-homogeneous classes would often neighbor each other in projections while classes with patterns were isolated individually along the outer border of the projection. The TCNN clustered classes with higher semantic homogeneity compared to the WCNN which clustered classes with similar frequency features together. Among both models, classification scores were satisfactory at roughly 98% for the TCNN and 95% for

the WCNN which supports why feature spaces segmented well.

5.3.2 ALOT

The analysis of ALOT demonstrated how strong texture overlap among its 250 classes can lower the precision of clustering individual classes but improve the class connectivity in feature space which can provide insight into the topological arrangement of features in the texture representation. Among the TCNN and WCNN experiments, both achieved lower classification scores around 74% and produced nearly identical feature space projections; however, the clustering varied in that the WCNN did not split classes containing larger objects and features as well because its frequency analysis components are not tailored for capturing shape information.

5.3.3 LPBF Datasets

Among the LPBF sensor modalities, all of them demonstrated high textural overlap as nearly all samples grouped into a single cluster in their projections similar to the ALOT projections. As a result, we focused on discussing region semantics and transitions between semantic categories. In following sections, we summarize our findings for each sensor modality.

5.3.4 Recoat Post-Spread

Segmentation

Powder spreads clustered into smooth, fine, and coarse categories. Other phenomenon including soot, minor recoat streaking, spotted and corner material protrusion grouped into their own clusters. Within the projections, several cluster regions composed small variations in uniform powder spreads. A large sparse region in each projection was dedicated

to powder spreads experiencing material protrusion. Among the sparse region, there existed a small cluster primarily capturing severe anomalies including recoat streaking, large material protrusion and non-uniform powder spreads.

Challenges

There was one critical problem with this dataset. The texture distribution from each coupon was unique; therefore, positioning within the build plane often determined the texture of the powder spreads. This was realized by several large regions in the projection and clusters grouping samples from a single coupon together because their average powder spread features were similar. These clusters often could not distinguish samples by the presence of soot/spatter. We speculate this position-texture bias was induced by both variations in lighting reflections across the build plane and blur from a misalignment between the build plane and camera focal plane. Both influences change the low-level pixel values which are embedded by the first feature maps in the TCNN and the first wavelet layer in the WCNN. Other larger-scale influential factors include dead pixel residue lines left over from mitigating sensor damage on coupons 4 and 5 as well as the distribution of soot being more common in back row coupons due to more turbulence towards the back of the argon air-flow over the build plane. Another challenge we faced was the distinction between infill material protrusion and corner protrusion; however distinguishing them becomes easy if the part geometry is included in the analysis.

Experiments

While Kylberg-TCNN produced the most clusters with one capturing bright speckled material protrusion, the Kylberg-WCNN produced a cluster capturing most of the severe defects with some false positives. The ALOT-TCNN clustering isolated coarse soot cleanly, while the ALOT-WCNN experiment had no notable additions.

Summary

In summary, we could effectively isolate powder spreads by roughness and the presence of corner protrusion but most soot and severe defects could not be clustered cleanly. The distributions of roughness levels in powder spreads heavily depended on the coupon positioning in the build plane thereby degrading the applicability of this approach and clusters to powder spreads unless corrective modifications to the data collection and preprocessing are made to equalize the sharpness and lighting angles across the build plane.

5.3.5 Recoat Post-Melt

Segmentation

The analysis of the recoat post-melt data revealed melted surfaces could be clustered by roughness levels including coarse, fine, rough, smooth, glossy, and saturated as well as by subtle features including etching and grooves. Within the projections, these roughness categories were arranged in a linear progression from coarse to rough, smooth, glossy, and finally saturated. The fine category would often reside alongside the transition from coarse to rough regions.

Challenges

Similar issues in the recoat post-spread modality came up in this modality too which we contribute to the fact that both visible recoat modalities use the same camera. Coupon positioning was still an issue; however, its influence did not degrade the utility of the approach as much for two reasons.

First, the texture distribution from each coupon spanned most roughness categories more uniformly, which is desirable. Samples grouped together by coupon row instead of by individual coupon. Nearly all fine melted surfaces came from the middle row coupons while

over-saturated coarse surfaces primarily belonged to front row coupons. We contribute these biases more towards the camera focus variations mentioned in the recoat post-spread conclusions than the lighting conditions because the camera focus varied most across the row axis in the build plane while the lighting produced a large gradient across the column axis in the build plane.

Second, the dead pixel residue lines blended in with melted surfaces even after normalization making it less likely for them to be detected and embedded into the texture representation. The reason for this was due to the fact that melted material often appears darker than its powder spread which brings its average pixel value down closer to the darker pixel values of the dead pixel residue lines. Additionally, the darkest pixels from shadows cast by bumps in the melted surfaces were typically darker than the dead pixel residue lines thereby making the residue pixels stand out less in comparison. For these reasons, we believe dead pixel residue lines were not factor in the segmentation of the feature space.

Experiments

The Kylberg-WCNN produced the semantically finest clustering while the other experiments often intermingled coarse and rough surface categories in larger clusters. The recoat post-melt was the only modality where the Kylberg-TCNN produced the fewest clusters. ALOT experiments had no notable benefits.

Summary

In summary, we could characterize melted surfaces by roughness categories with high accuracy; however the texture distributions heavily depended on the row positioning in the build plane thereby degrading the utility of the approach unless similar modifications mentioned in the recoat post-spread analysis are made.

5.3.6 Thermal Tomography

Segmentation

The analysis of the thermal tomography data revealed thermograms could be clustered by roughness levels including fine, coarse, rough, semi-rough, smooth, and saturated. Additionally, other features such as grains, central reliefs, large gradients, and spatter could be isolated into their own clusters except for cold spots which were absent from this dataset. The feature spaces revealed a smooth regional transition from fine to coarse, semi-rough, rough, smooth, and finally saturated.

Challenges

Thermal tomography only experienced a minor imbalance in the texture distribution across the rows in the build plane. Back row coupons were not as fine. Middle row coupons were not as rough. We suspect these biases might be related to the camera focus producing sharper imagery around the center row in the build plane. Additionally, coarse thermograms from front row coupons featured more bright/dark edges which may have been caused by imprecise image cropping.

Experiments

The Kylberg-TCNN produced the finest clustering, but could not isolate spatter into its own cluster; however, this was achieved for some rough and smooth instances in the ALOT-TCNN experiment. In general, both ALOT experiments were less effective at separating coarse and rough thermograms in their clusterings, but better at grouping spatter in their projections.

Summary

We could cleanly isolate thermograms by roughness category, grains, gradients, reliefs, and spatter. Modifications to the camera focus and image cropping that eliminate texture bias in future data are not immediately necessary but should eventually be implemented.

5.3.7 LWIR

Segmentation

LWIR heat maps clustered into crisp, coarse, rough, blurry rough, smooth, and blurry smooth categories. Less common features such as dark central reliefs and gradients formed separate clusters but spatter did not. The projections demonstrated there was a smooth transition from crisp to coarse, rough, and finally smooth categories in each feature space. Within the rough and smooth regions, there was a transition to blurry rough and blurry smooth sub-regions.

Challenges

There were two primary issues that degraded the applicability of the texture embedding. First, peak temperature heat maps produced subtle dark lines within the hatch spacing between laser passes. These lines were magnified after normalization. Although they varied in subtlety, every heat map captured these lines. We believe spatter was more difficult to cluster because these lines interfered with its embedding.

Second, coupon positioning had a strong influence on the texture distribution. Unlike in the recoat post-melt and thermal tomography modalities, texture distributions were unique to the column instead of the row in the build plane. Heat maps were blurry if they came from coupons in the left build plane column. Similarly, heat maps were crisp if they came from coupons in the right build plane column. We suspect this was a result of non-uniformity in

the camera focus on the build plane. Assuming it is a camera focusing problem, the reason it occurs over columns instead of rows is because the azimuth of LWIR camera relative to the build plane is rotated 90 degrees compared to the visible recoat and thermal tomography cameras.

Experiments

The Kylberg-WCNN clustering performed the best at segmenting roughness categories while the ALOT-WCNN clustering performed the worst grouping nearly all the data into a single cluster. Across all experiments, the topology of the projection were similar. The ALOT-TCNN projection was best at isolating a pure region of spatter.

Summary

We could effectively segment heat maps by roughness categories, dark central reliefs, and gradients. Additionally, spatter could easily be labeled in the projections. Before this approach can become applicable, the data collection and preprocessing require modifications that address the non-uniform blurring across the build plane and presence of laser scan lines in the imagery.

5.3.8 Spatter

Segmentation

Spatter emission maps clustered into rough, semi-rough, and smooth categories. Less common phenomena including dark central reliefs, small and large lattices, and bright edges grouped into separate clusters but spatter did not. The projections demonstrated a transition from rough to smooth regions within the feature spaces. Among these two major regions, large lattices gathered in the rough region while central reliefs and bright edges gathered in the smooth region.

Challenges

There were two issues with the data that may have degraded the utility of this approach. First, the presence of lattice artifacts was introduced by the data collection and preprocessing. The reason is related to overlap between the periodicity of the camera frame rate and the laser hatching strategy. These lattices varied in subtlety and size. Second, the texture distribution was slightly different between coupons in different columns. The primary observation was emission maps from right column coupons were slightly sharper at the pixel level. This implied there may be a subtle blur on the other coupons.

Experiments

The Kylberg-TCNN produced the finest clustering. Other experiments isolated less common phenomena better at the expense of poorer clustering among roughness categories. Most notably, the ALOT-TCNN clustering grouped most of the data into a single cluster thereby mixing smooth and rough categories but cleanly grouped dark central reliefs and over-saturation into their own clusters.

Summary

We demonstrated spatter emission maps grouped together by roughness, central reliefs, lattices, and bright edges via clustering. Although spatter was not segmented in the clustering, it could easily be labeled in the projections. Modifications to the data collection and preprocessing to remove lattices and equalize blurring are not required but would be beneficial in future research and applications.

5.3.9 Model Evaluation

From the analysis of all these datasets, we observed a few trends in each experimental model. The Kylberg-TCNN more often split datasets into the most number of clusters

which yielded the cleanest overall separation. The Kylberg-WCNN often competed with the Kylberg-TCNN sometimes producing the better clustering. Among both Kylberg experiments, feature spaces and clusters were overall similar. The ALOT-TCNN clusterings were competitive with those from Kylberg experiments; however the inability to distinguish rough samples from non-rough samples was a common issue. One benefit of the ALOT-TCNN experiment was clearer separation of spatter in thermal tomography clusters and the projections of LWIR and spatter feature spaces. The ALOT-WCNN usually provided the worst clusters often mixing roughness categories and showing no additional benefits compared to other experiments. Overall, the ALOT experiments performed worse at separating roughness categories but better at separating small round phenomenon such as spatter and soot compared to the Kylberg experiments. Part of the issue with ALOT experiments lacking good separation in roughness categories may be related to the difficulty of the ALOT classification problem due to having a magnitude more classes. We suspect the reason rarer phenomenon were better isolated is related to ALOT-specific classes providing the necessary features to learn features necessary for that segmentation.

5.4 Final Thoughts

We conclude with a few thoughts. Our approach was not tailored to any specific LPBF sensor modality. Despite that, we demonstrated unsupervised clustering could segment LPBF data by several roughness categories including fine, coarse, rough, and smooth. In addition, anomalous features such as dark central reliefs, gradients, and bright corners/edges were isolated. Given these capabilities, we believe the Kylberg texture embeddings can be used for time-efficient labeling and prediction of whether these features are present in new samples. This may provide an immediate benefit to process phenomenon classification across many sensor modalities; however, our approach could still be improved as it faced difficulties in clustering process phenomena with a small footprints such as spatter particles.

Additionally, texture biases from data collection and artifacts from preprocessing need to be removed before this approach can truly be applicable. Modifications for improvements and new strategies are reserved for future work in [Chapter 6](#).

Future Work

6.1 Modifications

6.1.1 Data Collection and Preprocessing

How data is collected and preprocessed should be considered carefully. Data analytics can be degraded if the data collection process or preprocessing destroys insightful information or introduces artifacts. For this reason we discuss strategies that mitigate texture biases in the LPBF data.

Camera Focusing

Blurring induced by the target being out of focus can quickly reduce the granularity of the image and change the overall texture. This is especially problematic when the blur is not equally distributed across the focal plane. This occurred significantly in the visible recoat and LWIR modalities and partially in the thermal tomography and spatter modalities. The severe instances demand correction before accurate data collection and characterization can be realized. To address this problem, we recommend installing a calibration lens onto each sensor to ensure regions in the build plane all have equal number of pixels in corresponding regions in the captured images. Additionally, when performing camera calibration with homography, the calibrated image resolution should be tailored towards the side of the

build plane with the fewest pixels to prevent upscaling in the calibration step from inducing a texture bias across the build plane.

Lighting

Lighting often varies in any image. The lighting orientation, reflections, and shadows may produce features in an image that dominate features of interest. This was primarily an issue for the visible recoat modalities. In the post-spread data, the lighting was not uniform across the column axis of the build plane. In the post-melt data, coupon reflections would vary based on build plane positioning and the processing parameters. While normalization was used to negate variations in the average pixel brightness, lighting conditions induced unique low-level patterns in powder spreads and reflections in melted surfaces across the build plane. Modifications to the physical arrangement of machine components are limited. Placing a uniform surface light above the build plane would obstruct the laser beam and the view of the sensors. A non-intrusive approach should be taken. For the post-spread data, we recommend implementing a lighting calibration preprocessing step. By capturing local pixel distributions across the first powder spread, subsequent powder spread pixels could be calibrated locally using these distributions. For the post-melt data, we have not found any viable strategies to discuss at this time.

Sensor Damage

The visible recoat camera had prior laser damage which yielded a thick line of dead pixels in the captured the image. These pixels were black and represent missing data; however, they also darkened their neighboring pixels. While we could have replaced both dead and darkened pixels with clean neighbors, the amount of replacement would start creating an artifact of its own. For this reason, we used the darkened pixels to replace the dead ones as our Goldilocks zone; however, these pixels were still embedded into the texture representation which caused artificial separation between coupons 4 and 5 in the powder

spreads. While there are other image filling techniques that may calibrate the darkened pixels to their correct brightness, we recommend replacing the sensor as it is inexpensive.

Normalization

In prior experiments that did not individually normalize images, we found textures segmented more often by average brightness than roughness and anomalies. For that reason we started using normalization to standardize the mean and variance per image sample to ensure texture was prioritized. Although this preprocessing performed well, it did have some side effects. When very bright or dark corners/edges/spots deviating from the mean pixel value were present, they raised the standard deviation of the image which resulted in the other pixel values being squashed down by normalization. This led to textures being smoother if they featured bright or dark pixels that deviated far from the mean. We recommend replacing normalization with a mean-shift transform to ensure images have a zero mean without changing their variance.

Compiling Frames

In thermal tomography, LWIR, and spatter, multiple frames were captured per layer. To obtain a single texture image for the layer, these frames had to be compiled together; however the method of compilation yielded artifacts in LWIR and spatter. Max temperature LWIR heat maps featured bright lines from the laser pass and darker lines within the hatch spacing. Almost all image samples exhibited these lines. Some spatter maps embedded varying lattice structures where pixels on these lattices were brighter. In both modalities, these artifacts might skew analysis using new data that does not feature these artifacts.

Cropping

In this work, images were cropped to fit the square coupon regions. This was done to ensure 100% of the image modeled a texture pixel distribution that excluded sharp edges at the coupon boundary. By cropping in this way, any part geometry can be analyzed by convolving a window within the infill region of the part and embedding those regions into a texture representation that can be characterized by finding a nearest neighbor in a previous clustering; however this may not hold if boundary information is captured. It is for this reason that we recommend tightening the cropping as some samples from front row coupons in thermal tomography captured dark edges that over-extended beyond the coupon boundary.

6.1.2 Model Architecture

Energy Layer

In the TextureCNN we used average pooling energy layers on feature maps from each level to embed both low-level and high-level features into the latent vector. There were two issues with this configuration. First, textures in the LPBF data were not semantically uniform within the image window when anomalies were present. Most anomalies had a small footprint within the coupon, so their filter responses would be washed out by average pooling in the energy layers because the rest of the feature map would have low responses for those filters where the anomaly was not present. We suggest supplementing the TextureCNN feature maps with max pool energy layers to ensure small, uncommon features are included in the embedding. Second, low-level features were embedded into the texture space which helped characterize roughness categories; however due to blurring, this could promote the model's sensitivity to low-level texture distribution biases. This event was most extreme in the visible recoat post-spread images. We recommend removing the energy layer from the first set of feature maps to restrain pixel-level information from being prioritized.

Bottleneck Layer

The bottleneck layer was responsible for embedding features into a compact latent vector that could be used to classify the input via logistic regression in the last fully-connected layer. This bottleneck layer was activated using a ReLU function; however, a few of its neurons died during training. This means not every dimension in the latent space was used. We recommend replacing ReLU with leaky ReLU or hyperbolic tangent as their gradients are non-zero for their entire domain. This should prevent neurons from dying and ensure all latent dimensions are utilized thereby producing a less entangled and more complete texture representation.

6.1.3 Clustering

The LPBF feature space projections demonstrated there was high textural overlap within the images from each sensor modality. This overlap led to distinct roughness categories in the feature space being connected by hundreds of samples providing a smooth interpolation between the two categories. As a result, the cut between clusters assigned by FDC (hard clustering) became more arbitrary. For this reason, we recommend other clustering approaches that describe the roughness labeling of clusters in a continuous domain rather than a discrete one. Soft and hierarchical clustering algorithms such as Fuzzy DBSCAN [13] and HDBSCAN [20] may be better suited for this task.

6.2 LPBF Tailored Approach

The strategy presented in this thesis relied on learning a compact texture representation from features present in general texture datasets; however, each LPBF sensor modality and its corresponding preprocessing produced a unique scope of texture images. If the learned texture representation was scoped specific to each sensor modality, cleaner segmentation

with purer corresponding semantics may be achieved more easily. In deep learning literature, there exist architectures that can do this for unlabeled datasets such as the LPBF data. These architectures build off a basic model, namely an autoencoder [10]. The task of an autoencoder is to learn a compact representation by embedding an input into a low-dimensional (latent) space such that its reconstructed output matches the input. In recent literature, this basic model has been tailored to various tasks by supplementing it with other mechanisms that improve performance. One of these developed architectures is a Variational Autoencoder (VAE) [14]. It changes how inputs are embedding by reparameterizing them into mean and standard deviation components. This helps the model generalize feature combinations for input reconstructions and improve the connectivity and explorability of the latent space. Training a VAE on a dataset can produce a latent space that better isolates and organizes semantic information. By sampling the latent space with reconstructions, clustering and labeling become more effective.

The utility of a VAE does not end with semantic labeling and characterization. In another material manufacturing domain, a TextureVAE [1] was developed using a pre-trained VGG19 network and a style loss function to learn a representation for material microstructure images (i.e. cast iron). The latent space could be traversed along a single dimension to vary a single attribute of generated microstructure images. Additionally, mechanical properties tied to microstructure elements were highly predictive in the latent space embedding. This was demonstrated by training a linear regression model to map the dimensions of the latent space to the mechanical properties in which it achieved R^2 scores just shy of 0.9. This work is very relevant to the process development challenges faced in LPBF. To demonstrate, we propose a series of steps that could be performed for each sensor modality to map process parameters to process monitoring imagery.

1. Train a TextureVAE to learn a textural representation for a wide variety of sensor data.
2. Cluster and label embedded images in the latent space.

3. Train simple ML models to map latent dimensions back to processing parameters.
4. Explore the processing parameter space by visualizing generated sensor images using the learned mapping between parameters and the latent dimensions.
5. Find processing parameters that produce desirable process monitoring imagery.

Given the potential of texture embedding, it is crucial that the strategies demonstrated in this approach be applied in further research to improve process quality control and parameter development with process monitoring. We initially attempted a similar solution by training a smaller TextureVAE from scratch, but learning to generate sharp reconstructions was an obstacle we did not overcome at the time. This would be a project to revisit.

Bibliography

- [1] Texturevae : Learning interpretable representations of material microstructures using variational autoencoders. 2021.
- [2] AddiGuru. Product.
- [3] alexandreday. Fast density clustering. https://github.com/alexandreday/fast_density_clustering, 2018.
- [4] Vincent Andrearczyk and Paul F. Whelan. Using filter banks in convolutional neural networks for texture classification. *CoRR*, abs/1601.02919, 2016.
- [5] Hermann Baumgartl, Josef Tomas, Ricardo Buettner, and Markus Merkel. A deep learning-based model for defect detection in laser-powder bed fusion using in-situ thermographic monitoring, Feb 2020.
- [6] G. J. Burghouts and J. M. Geusebroek. Material-specific adaptation of color invariant features. *Pattern Recognition Letters*, 30:306–313, 2009.
- [7] Andy Coenen and Adam Pearce. Understanding umap, 2019.
- [8] Shin Fujieda, Kohei Takayama, and Toshiya Hachisuka. Wavelet convolutional neural networks for texture classification, 2017.

- [9] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS' 15, page 262–270, Cambridge, MA, USA, 2015. MIT Press.
- [10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [11] Marco Grasso, Afaf Remani, Andrew Dickins, Bianca Colosimo, and Richard Leach. In-situ measurement and monitoring methods for metal powder bed fusion – an updated review. *Measurement Science and Technology*, 32, 05 2021.
- [12] R.M. Haralick, K. Shanmugam, and I Dinstein. Textural features for image classification. *IEEE Transactions on In Systems*, 1973.
- [13] Dino Ienco and Gloria Bordogna. Fuzzy extensions of the dbscan clustering algorithm. *Soft Comput.*, 22(5):1719–1730, mar 2018.
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014.
- [15] Tomasz Kurzynowski, Wojciech Stopyra, Konrad Gruber, Grzegorz Ziółkowski, Bogumiła Kuźnicka, and Edward Chlebus. Effect of scanning and support strategies on relative density of slm-ed h13 steel in relation to specimen size. *Materials*, 12:239, 01 2019.
- [16] G. Kylberg. The kylberg texture dataset v. 1.0.
- [17] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010.
- [18] L. McInnes and J. Healy. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *ArXiv e-prints*, February 2018.
- [19] Leland McInnes. Using umap for clustering, 2018.

- [20] Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2, 03 2017.
- [21] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap, 2018.
- [22] Keiron O’Shea and Ryan Nash. An introduction to convolutional neural networks, 2015.
- [23] C. Panwisawas, C.L. Qiu, Y. Sovani, J.W. Brooks, M.M. Attallah, and H.C. Basoalto. On the role of thermal fluid dynamics into the evolution of porosity during selective laser melting. *Scripta Materialia*, 105:14–17, 2015.
- [24] 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.
- [25] Luke Scime and Jack Beuth. Anomaly detection and classification in a laser powder bed additive manufacturing process using a trained computer vision algorithm. *Additive Manufacturing*, 19:114–126, 2018.
- [26] Luke Scime, Derek Siddel, Seth Baird, and Vincent Paquit. Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation. *Additive Manufacturing*, 36:101453, 2020.
- [27] M. Unser. Texture classification and segmentation using wavelet frames. *IEEE Transactions on Image Processing*, 4(11):1549–1560, 1995.
- [28] Manik Varma and Andrew Zisserman. A statistical approach to texture classification from single images. *Int. J. Comput. Vision*, 62(1–2):61–81, April 2005.

- [29] Daniel Vriesman, Alessandro Zimmer, Alceu S. Britto Jr. au2, and Alessandro L. Koerich. Texture cnn for thermoelectric metal pipe image classification, 2019.
- [30] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. *CoRR*, abs/1604.03628, 2016.
- [31] Dongxiao Zhou. Texture analysis and synthesis using a generic markov-gibbs image model, 2006.