

DECOR: Deep Embedding Clustering with Orientation Robustness

Fiona Victoria Stanley Jothiraj^{1*}, Arunaggiri Pandian Karunanidhi², Seth A. Eichmeyer²

¹Oregon State University

²Micron Technology

stanleyf@oregonstate.edu, akarunanidhi@micron.com, saeichmeyer@micron.com

Abstract

In semiconductor manufacturing, early detection of wafer defects is critical for product yield optimization. However, raw wafer data from wafer quality tests are often complex, unlabeled, imbalanced and can contain multiple defects on a single wafer, making it crucial to design clustering methods that remain reliable under such imperfect data conditions. We introduce DECOR, a *deep clustering with orientation robustness* framework that groups complex defect patterns from wafer maps into consistent clusters. We evaluate our method on the open source MixedWM38 dataset, demonstrating its ability to discover clusters without manual tuning. DECOR explicitly accounts for orientation variations in wafer maps, ensuring that spatially similar defects are consistently clustered regardless of its rotation or alignment. Experiments indicate that our method outperforms existing clustering baseline methods, thus providing a reliable and scalable solution in automated visual inspection systems.

1 Introduction

Clustering is an unsupervised learning technique used to group unlabeled data based on the underlying similarity in their feature representations. Methods are often categorized into hard clustering and soft clustering based on the exclusivity of clusters: hard clustering assigns each data point to exactly one cluster, whereas soft clustering allows a data point to belong to multiple clusters with varying probabilities. When dealing with high-dimensional data such as wafer maps, autoencoder variants are commonly used to extract the low-dimensional latent embeddings that preserve structural information. These embeddings (often a 1-dimensional vector) are subsequently clustered using well-established approaches, including centroid-based clustering (Eg. K-Means), density-based clustering (Eg. DBSCAN, HDBSCAN, OPTICS), distribution-based clustering (Eg. Gaussian Mixture Models, Expectation-Maximization) or hierarchical clustering (Eg. Agglomerative Clustering).

Despite their effectiveness and widespread use, most clustering methods rely on fixed assumptions, such as the number of clusters (k), maximum distance between points (ϵ),

and choice of distance metric (Eg. L1 or L2 distance). This reliance of assumptions makes it impractical for dynamic wafer inspection, where the wafer defect distribution can evolve over time, and clustering needs to be applied frequently to identify outlier or anomalous patterns without prior knowledge of their prevalence or scale.

Within semiconductor manufacturing, this challenge is compounded by the fact that wafer maps often exhibit orientation variability. The same defect pattern may appear at different rotations due to variations of wafer placement or handling. Standard clustering methods treat rotated instances as being distinct patterns and this leads to fragmented clusters.

These challenges motivate the adoption of non-parametric methods of clustering (Nguyen et al. 2021; Xing et al. 2021; McConville et al. 2021) in our work which can adaptively discover clusters in latent space without requiring prior specifications, while also aligning spatially similar wafer defects across orientations. We refer to our approach as DECOR (**D**eep **E**mbedding **C**lustering with **O**rientation **R**obustness).

2 Method

We propose a novel three-step approach (Algorithm 1) for non-parametric wafer map clustering and cluster-aware outlier detection (Figure 1). Our approach consists of: **(A)** a rotation- and flip-invariant embedding extractor - RCAE, **(B)** a non-parametric clustering module, and **(C)** an ensemble outlier detection mechanism.

2.1 Feature Extractor

To obtain robust and well-separated embeddings from high-dimensional wafer maps, we evaluate several architectures, including MoCo (He et al. 2019; Chen et al. 2020b,a) for contrastive representation learning and convolutional autoencoders (Masci et al. 2011) for reconstruction-based learning. While MoCo has shown effectiveness in general contrastive learning tasks, it lacks explicit reconstruction and its architecture does not inherently enforce symmetry handling. However, our experiments indicate that a rotation invariant convolutional autoencoder (CAE (Masci et al. 2011)) with built-in symmetry handling produced more compact and separable clusters in the reduced embedding space (ie. RCAE). We processed wafer maps of size 128×128 through three *rotation- and flip-equivariant* blocks, each comprising an R2Conv layer (Weiler and Cesa 2019),

*Work conducted as part of 2025 summer internship at Micron Technology.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Algorithm 1: DECOR: Dynamic Embedding Clustering with Orientation Robustness

Require: Images $\mathcal{I} = \{I_1, \dots, I_N\}$; embedding dimension d ; initial number of clusters K_{init}
Ensure: Soft memberships $\mathbf{P} \in [0, 1]^{N \times K}$; hard labels $\mathbf{c} \in \{1, \dots, K\}^N$; outlier set \mathcal{O}^*

```

1: for  $i = 1$  to  $N$  do
2:    $I_i \leftarrow \text{NORMALIZEANDMASK}(I_i)$ 
3:    $z_i \leftarrow \text{FEATUREEXTRACTOR}(I_i) \in \mathbb{R}^d$ 
4: end for
5:  $\mathbf{Z} \leftarrow [z_1; \dots; z_N] \in \mathbb{R}^{N \times d}$ 
6:  $\mathcal{A} \leftarrow \text{DYNAMICCLUSTERING}(K_{\text{init}})$ 
7:  $\mathbf{P} \leftarrow \mathcal{A}.\text{SOFTCLUSTER}(\mathbf{Z}) \in [0, 1]^{N \times K}$ 
8:  $\mathbf{c} \leftarrow \arg \max_{k \in \{1, \dots, K\}} \mathbf{P}$  ▷ Hard labels from soft memberships
9:  $\mathcal{O}^* \leftarrow \bigcup_{k=1}^K \text{OUTLIERDETECTION}(\{z_i \mid c_i = k\})$  ▷ Outlier detection per cluster
10: return  $\mathbf{c}, \mathcal{O}^*$ 

```

ReLU layer for non-linearity, and PointwiseAvgPool with stride 2. The blocks use increasing numbers of regular representation fields (8, 16, 32), achieving equivariance to the dihedral group D_4 - four discrete rotations ($0^\circ, 90^\circ, 180^\circ, 270^\circ$) and two mirror flips. To convert the equivariant features into orientation-invariant descriptors, we applied GroupPooling, which collapses all orientation channels within each field. A Linear layer then maps the resulting tensor to a 128-dimensional latent vector. The decoder mirrors the encoder, using three ConvTranspose2D+ReLU blocks and a final Sigmoid activation to reconstruct images within $[0, 1]$.

2.2 Non-Parametric Clustering

For clustering in the latent space, we use DeepDPM (Ronan, Finder, and Freifeld 2022), a non-parametric deep clustering framework based on the Dirichlet Process Mixture Model (DPMM). Unlike parametric clustering approaches like K-Means, which require a fixed number of clusters, DeepDPM infers the optimal number of clusters from the data distribution. The trained clustering model is implemented as a two-layer multilayer perceptron comprising a Linear(128, 50) layer followed by Linear(50, K), where K is adaptively determined during training. The model produces soft cluster membership probabilities, from which hard assignments are obtained.

2.3 Outlier Detection

In the latent space produced by the invariant encoder, outlier detection is performed with an ensemble of Isolation Forest (IF) (Liu, Ting, and Zhou 2008, 2012) and Local Outlier Factor (LOF) (Breunig et al. 2000). IF provides a global, partition-based notion of isolation, while LOF captures local density deviations. Their raw scores are converted to outlier decisions using robust, per-cluster thresholding rule. For the score vector s , we define the decision threshold as $\tau = \text{median}(s) + k \cdot \text{MAD}(s)$ where the Median Absolute Deviation (MAD) is given by $\text{MAD}(s) = \text{median}(|s - \text{median}(s)|)$. We use the median and MAD rather than mean and standard deviation to ensure robustness against heavy-tailed distributions and pre-existing outliers. To stabilize LOF across heterogeneous cluster sizes, the neighborhood size was chosen adaptively as $k_{\text{LOF}} =$

$\text{clip}(\sqrt{N}, k_{\text{min}}, k_{\text{max}})$ for cluster cardinality N . IF was trained with a conservative contamination prior ($\text{hi_cont} = 0.20$) to avoid over-flagging, while the subsequent robust cut calibrates per-cluster selectivity. Final outlier labels were issued only when both detectors agree, $\mathbf{1}_{\text{final}} = \mathbf{1}_{\text{IF}} \wedge \mathbf{1}_{\text{LOF}}$. This process empirically improves precision on rare, semantically novel wafer defects while retaining sensitivity through the complementary failure modes of the two detectors.

3 Experiment

3.1 Dataset

We train and evaluate our method on the opensource MixedWM38 (Wang et al. 2020) dataset, which contains over 38,000 wafer maps with both single-defect and mixed-defect patterns, spanning more than 38 distinct defect pattern combinations. In particular, the dataset includes 1 normal pattern, 8 single-defect and 29 mixed-defect patterns. The base defect types are *Center*, *Donut*, *Edge-Loc*, *Edge-Ring*, *Loc*, *Near-full*, *Scratch*, and *Random*. To ensure consistent multi-label distribution across training and testing, we utilize MultilabelStratifiedShuffleSplit (Sechidis, Tsoumakas, and Vlahavas 2011), which preserves the proportion of each defect type in both subsets. Additionally, all wafer maps are resized to 128×128 pixels to balance reconstruction fidelity and clustering performance.

3.2 Results

We evaluate the clustering quality using standard metrics like: *Normalized Mutual Information (NMI)* and *Adjusted Rand Index (ARI)*. Since the ground truth (GT) is multi-label (i.e. single wafer can contain multiple defects), we perform a cluster-aware dominant label reduction of GT: for each wafer, if multiple defect labels are present, the most frequent within the sample's predicted cluster is assigned. NMI and ARI are then computed between these dominant labels and the predicted cluster assignment. For robustness, we repeat experiments with multiple random seeds and report mean values and standard error across three runs.

A key challenge with this dataset is the absence of a well-defined ground-truth number of categories since wafers may

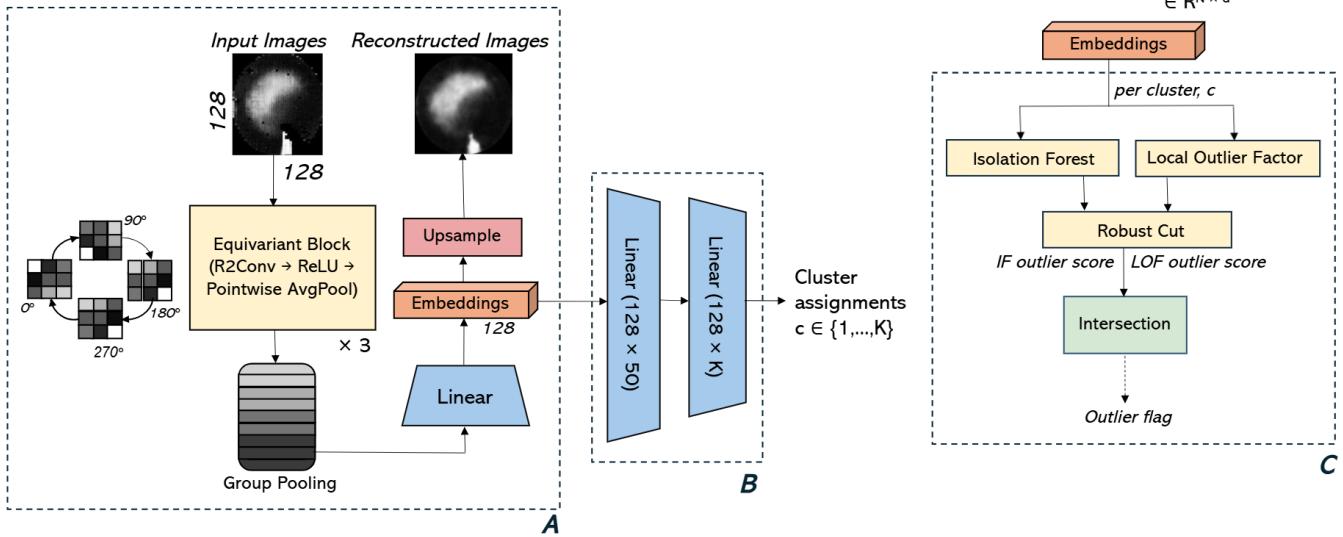


Figure 1: Overview of the workflow. Input wafer images are processed through an embedding extractor model, generating 128-dimensional embeddings along with their reconstructed images. These embeddings are then fed into a trained non-parametric clustering model based on DeepDPM (Ronan, Finder, and Freifeld 2022) to obtain cluster assignments. Finally, embeddings within each cluster are analyzed using an ensemble outlier detection algorithm to flag potential outliers.

contain overlapping or mixed defect types. This makes comparisons with parametric baselines (eg. K -Means) difficult, as their performance depends heavily on the choice of K .

Table 1 reports the performance of different embedding-clustering combinations. Our results indicate that RCAE embeddings, when paired with a non-parametric clustering approach (DeepDPM), achieves highest performance in both NMI and ARI. We observe that RCAE+DeepDPM produces more compact and orientation-invariant clusters compared to baselines (Figure 2). Furthermore, we find the quality of clustering correlates with the degree of separation between defect patterns in latent space: the more distinct the separation in 2D/3D projections (via (McInnes, Healy, and Melville 2018) (Maaten and Hinton 2008)), the better the clustering performance.

4 Conclusion

In this work, we address the challenges of reliable wafer map clustering under complex, noisy and multi-defect conditions. We propose DECOR, a novel framework that combines an orientation-invariant embedding extractor, a non-parametric clustering model and an outlier detection mechanism to group wafer maps into meaningful clusters, while flagging outlier defects. Experiments on the multi-label MixedWM38 dataset demonstrate that our method achieves superior clustering performance with minimal computational overhead, by leveraging lightweight yet effective models. One limitation of the non-parametric clustering approach is the necessity of multiple runs to determine an appropriate initial cluster count (k_{init}) and the optimal number of training epochs for DeepDPM clustering that is tailored to the dataset. Furthermore, evaluating clustering qual-

ity on multi-label datasets remains challenging, as metrics like NMI and ARI require simplification to dominant labels per image, which might not capture the full system's performance.

As part of future work, we plan to explore multi-label-aware clustering metrics, and integrate temporal modeling, leveraging soft cluster assignments produced by non-parametric clustering to track the evolution of defect patterns across production cycles. We also plan to benchmark DECOR against additional parametric clustering models, including Gaussian Mixture Models and deep clustering frameworks beyond K -means, to further assess its robustness and generalization. By advancing reliable unsupervised learning for imperfect and complex data, DECOR provides a foundation that generalizes beyond semiconductor wafers to other domains/datasets.

A Experimental Setup

We evaluate our method on the MixedWM38 wafer dataset using a multilabel stratified train-test split. After normalization, we applied Edge masking to prevent edge artifacts from dominating the embeddings. We also applied Gaussian blur transformation (kernel size 5, and sigma 1.0) to smooth noise and highlight the main defect patterns for clustering. Our embedding feature extractor was compared against baselines including vanilla CAE and MoCo, all trained under identical conditions for fair comparison. The RCAE model was trained using the Adam optimizer (initial learning rate: $1e^{-3}$), MSE loss, batch size of 128, and 1000 epochs. For DeepDPM clustering, we set hyper parameters as follows: ν is $d + 2$ (where d is the input dimension to the clustering model or the latent dimension), k_{init} is 30,

Table 1: Performance under different experiments indicating the mean \pm std. dev across 3 runs. Methods marked with p are parametric (require K). Experiment details are provided in *Appendix A-C*

Embedding	Clustering	Final/Best K	NMI \uparrow	ARI \uparrow
CAE	K -Means p	24	0.503 ± 0.00	0.199 ± 0.00
MoCo	K -Means p	18	0.409 ± 0.00	0.173 ± 0.01
RCAE	K -Means p	30	0.529 ± 0.00	0.205 ± 0.00
CAE	DeepDPM	25	0.498 ± 0.05	0.218 ± 0.01
MoCo	DeepDPM	30	0.273 ± 0.00	0.117 ± 0.00
RCAE (Ours)	DeepDPM	22	0.543 ± 0.03	0.296 ± 0.00

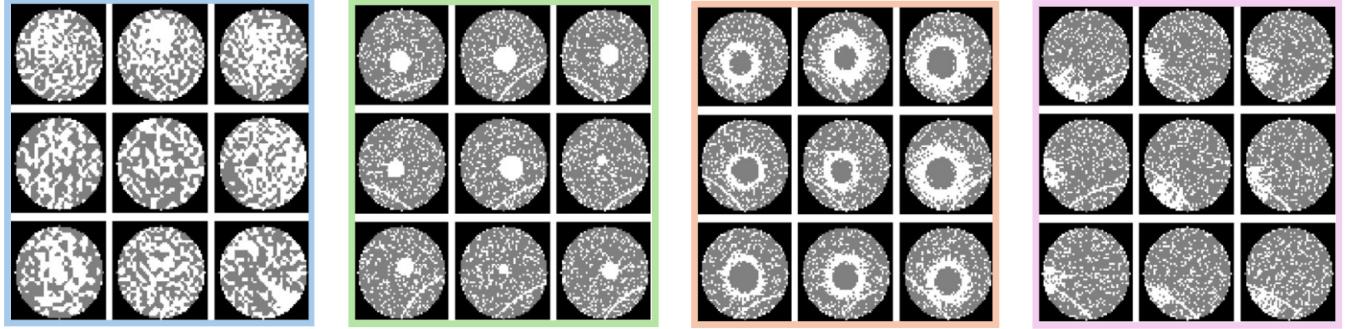


Figure 2: Examples of MixedWM38 images clustered together by DECOR. Each panel represents a different cluster. The grouping of images with varying rotational orientations (*scratch*, *and local* pattern) indicates that DECOR exhibits rotational invariance.

and maximum epochs is 200. Although labels were withheld during training, we report evaluation metrics using NMI and ARI.

All experiments were implemented in PyTorch 2.6.0 and run on a single NVIDIA H100 GPU (80GB memory) with 32-core CPUs on a local compute cluster node. Training the RCAE model for 1000 epochs took approximately 6 hours per run, while training the non-parametric clustering model required about 2 hours. Across all experiments (including baselines), the total compute amounted to approximately 30GPU hours. Storage requirements were minimal, with about 400MB for datasets and about 1GB for checkpoints and logs.

B Outlier Detection

B.1 Isolation Forest (IF) and Local Outlier Factor (LOF)

Outlier detection is an essential task in many domains, especially in wafer map analysis where it is crucial to distinguish between genuine defects and noise. Isolation Forest (IF) (Liu, Ting, and Zhou 2008, 2012) and Local Outlier Factor (LOF) (Breunig et al. 2000) are both widely used for outlier detection due to their ability to handle high-dimensional and complex data. However, each has its strengths and limitations, which is why an ensemble approach is adopted to leverage the advantages of both.

Isolation Forest operates by randomly partitioning the feature space and isolating data points. Outliers, being different

from normal points, are easier to isolate and therefore receive a lower score. On the other hand, Local Outlier Factor (LOF) focuses on local densities, comparing the density of a point with its neighbors. Points with significantly lower density than their neighbors are flagged as outliers.

While these algorithms are powerful on their own, they tend to excel in different types of data and outlier characteristics. To take advantage of both, an ensemble approach is employed, where outliers are flagged only if both Isolation Forest and LOF agree on the outlier, reducing the chance of false positives and increasing the robustness of the results.

B.2 Robust Cut and Thresholding

The performance of outlier detection algorithms depends heavily on defining appropriate thresholds for outlier scores. In this study, the Robust Cut approach is used to define these thresholds in a way that adapts to the data's underlying distribution. The Robust Cut uses the Median Absolute Deviation (MAD) to ensure that the threshold is not influenced by extreme values (outliers) in the score distribution.

The formula for the robust cut is defined as:

$$\text{Threshold} = \text{Median}(\text{scores}) + k \times \text{MAD}(\text{scores})$$

where $\text{Median}(\text{scores})$ is the median of the outlier scores, $\text{MAD}(\text{scores})$ is the median absolute deviation, calculated by $\text{MAD} = \text{Median}(|\text{scores} - \text{Median}(\text{scores})|)$, and k is a scaling factor, which controls how aggressive the thresholding is.

Table 2: Implementation details of feature extractor models.

Architecture	Loss	lr
CAE: Conv(16) - Conv(32) - Conv(64) - Conv(256) - FC(128)	MSE	0.001
MoCo: ResNet18 - FC(512) - FC(128)	InfoNCE	0.03
RCAE: R2Conv(8) - R2Conv(16) - R2Conv(32) - GroupPool - FC(128)	MSE	0.001

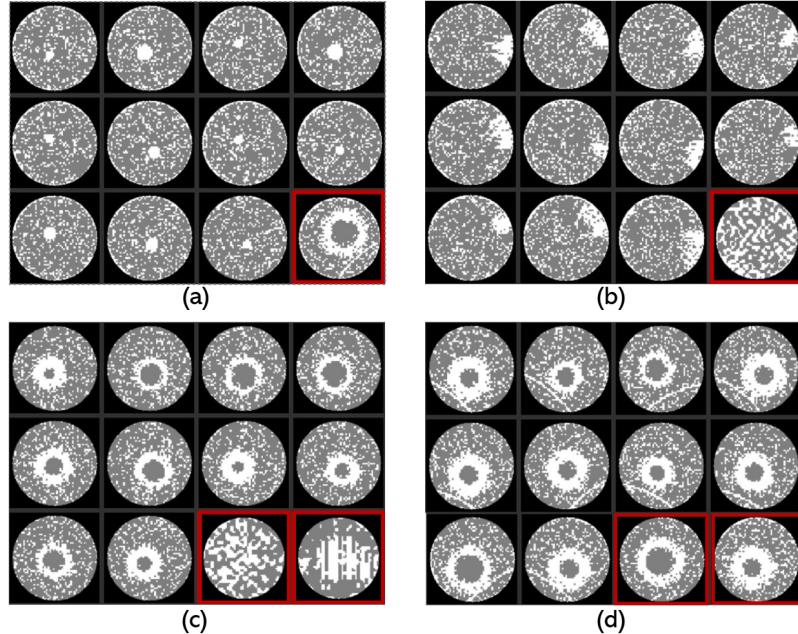


Figure 3: Examples of clusters produced by DECOR with detected outliers highlighted in red. (a) a *center* defect cluster containing instances of *donut* defect outliers (b) a *local* defect cluster containing instances of *random* defect outliers (c) a *donut* defect cluster containing instances of *random* and *nearfull* defect outliers (d) a *donut + scratch* defect cluster containing instances of *donut* defect outliers.

Median is preferred over the mean because it is less sensitive to extreme values, which are common in high-dimensional data. This ensures that the threshold is robust to noise or outliers in the data and avoids flagging too many points as outliers when the dataset contains a few extreme cases.

MAD provides a more robust estimate of spread, which is particularly important when dealing with outliers. By scaling the MAD with k , we can control how aggressively the algorithm flags outliers.

The contamination parameter (*hi_cont*) is set to 0.20, meaning that 20% of the data are expected to be outliers. This is a conservative choice to ensure that true outliers are not overlooked, especially in scenarios where outliers might be rare or subtle. Setting a low contamination value prevents the algorithm from being overly aggressive and helps in retaining the majority of the dataset as normal.

B.3 Adaptive Thresholding and LOF Parameters

The Adaptive k parameter in Local Outlier Factor (LOF) adjusts the number of neighbors considered when evaluating the density of each point. Instead of a fixed value

for k , we compute an adaptive k based on the size of the dataset - specifically the square root of the number of points in the cluster, constrained within a predefined range $[min_k, max_k]$. This adaptability ensures that the LOF algorithm remains sensitive to local structures regardless of cluster size, which is essential when dealing with clusters of varying densities.

The formula for the adaptive k is as follows:

$$k = \text{Clip}(\sqrt{N}, min_k, max_k)$$

where N is the number of data points in the cluster and $\text{Clip}(x, min_k, max_k)$ ensures that k remains within the bounds of $[min_k, max_k]$.

By choosing the appropriate k , LOF can accurately capture the local structure of the data and identify outliers that are dense in their neighborhoods but sparse relative to other areas.

B.4 Ensemble Outlier Detection

The ensemble approach combines the results of Isolation Forest (IF) and Local Outlier Factor (LOF) to determine whether an image is an outlier. The method is based on

the assumption that true outliers are detectable both globally (via IF) and locally (via LOF), so an intersection of both methods is used for a more reliable result. The final outlier decision is computed by the logical AND of the outlier flags from both algorithms: $\text{Final Outlier Flag} = \text{IF Flag} \cap \text{LOF Flag}$, where IF Flag is the result of applying Isolation Forest to the data, LOF Flag is the result of applying Local Outlier Factor to the data. This ensemble approach ensures that the outlier is detected by both methods before it is flagged as a true outlier, thereby reducing false positives and increasing the reliability of the outlier detection method.

C Clustering Results

Figure 3, shows examples of images from the MixedWM38 Validation set that were clustered using the RCAE feature extractor and DeepDPM non-parametric clustering approach. Outliers identified by the ensemble outlier-detection algorithm are highlighted with red outline. For outlier detection, we set the contamination prior parameters (hi_cont) of 0.20, robust cut scaling factor (k) of 3 for both IF and LOF algorithms. For the adaptive k parameter in LOF, we use a min_k and max_k of 10 and 50, respectively.

References

Breunig, M. M.; Kriegel, H.-P.; Ng, R. T.; and Sander, J. 2000. LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 93–104. New York, NY, USA: Association for Computing Machinery.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.

Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020b. Improved Baselines with Momentum Contrastive Learning. *arXiv preprint arXiv:2003.04297*.

He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2019. Momentum Contrast for Unsupervised Visual Representation Learning. *arXiv preprint arXiv:1911.05722*.

Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation forest. In *2008 eighth ieee international conference on data mining*, 413–422. IEEE.

Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2012. Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1): 1–39.

Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.

Masci, J.; Meier, U.; Cireşan, D.; and Schmidhuber, J. 2011. Stacked convolutional auto-encoders for hierarchical feature extraction. In *International conference on artificial neural networks*, 52–59. Springer.

McConvile, R.; Santos-Rodriguez, R.; Piechocki, R. J.; and Craddock, I. 2021. N2d:(not too) deep clustering via clustering the local manifold of an autoencoded embedding. In *2020 25th international conference on pattern recognition (ICPR)*, 5145–5152. IEEE.

McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.

Nguyen, X.-B.; Bui, D. T.; Duong, C. N.; Bui, T. D.; and Luu, K. 2021. Clusformer: A transformer based clustering approach to unsupervised large-scale face and visual landmark recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10847–10856.

Ronen, M.; Finder, S. E.; and Freifeld, O. 2022. DeepDPM: Deep Clustering With An Unknown Number of Clusters. In *Conference on Computer Vision and Pattern Recognition*.

Sechidis, K.; Tsoumakas, G.; and Vlahavas, I. 2011. On the stratification of multi-label data. In *Joint European conference on machine learning and knowledge discovery in databases*, 145–158. Springer.

Wang, J.; Xu, C.; Yang, Z.; Zhang, J.; and Li, X. 2020. Deformable convolutional networks for efficient mixed-type wafer defect pattern recognition. *IEEE Transactions on Semiconductor Manufacturing*, 33(4): 587–596.

Weiler, M.; and Cesa, G. 2019. General E(2)-Equivariant Steerable CNNs. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Xing, Y.; He, T.; Xiao, T.; Wang, Y.; Xiong, Y.; Xia, W.; Wipf, D.; Zhang, Z.; and Soatto, S. 2021. Learning hierarchical graph neural networks for image clustering. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3467–3477.