

Supplementary Material for FRE: A Fast Method For Anomaly Detection And Segmentation

Ibrahima J. Ndiour
ibrahima.j.ndiour@intel.com

Nilesh Ahuja
nilesh.ahuja@intel.com

Utku Genc
utku.genc@intel.com

Omesh Tickoo
omesh.tickoo@intel.com

Intel Labs
Intel Corporation
USA

We present here the supplementary material for the 2023 BMVC paper *FRE: A Fast Method For Anomaly Detection And Segmentation*. While in the main text, we sometimes offered class-average results for brevity and readability, we provide here the corresponding detailed class-itemized results. In addition, we also provide a few more results showing performance across various backbones and layers, as well as pictures showing performance for a few sampled input images.

Compute Performance: Table 1 shows the performance of our method and compares it against the state of the art on various computing platforms (laptop-class CPU, server-class CPU and a high-end discrete GPU).

Choice of Layers: Tables 2, 3, 4, and 5 respectively provide a class-average summary of anomaly detection and segmentation performance for FRE across various layers of EfficientNet_B5, Resnet18, Resnet50 and VGG16 backbones. In tables 8 and 9, we show the class-itemized segmentation performance of FRE across multiple layers of a Resnet50 backbone. We observe that the middle layers generally tend to provide better detection and location performance relative to layers closer to either the input or the output.

Additional Outputs: Figure 1 shows a few examples of segmentation maps obtained at various layers of a Resnet50 backbone, as well as the segmentation maps obtained with the combination of 3 layers. Figure 2 shows examples of the segmentation maps produced by FRE with a single layer, and with the combination.

Implementation strategies: Figure 3 shows the (class-itemized) empirical validation of the different implementation strategies presented in the main paper. It is seen that all strategies

converge to the same anomaly detection performance thereby demonstrating that the three strategies are basically equivalent.

Sensitivity to parametrization: Figures 4 and 5 provide the class-itemized detection and segmentation performance for FRE while varying the only parameter of the method, for multiple layers of a Resnet50 backbone. These figures not only show insensitivity to parametrization, but they also give a more detailed view of the performance of the method.

Choice of backbones: Tables 6 and 7 shows detailed class-itemized pixel-wise AUROC metrics for the segmentation benchmark and for FRE across backbones.

	Training Time (↓)				Inference frames-per-sec (↑)		
	PatchCore[0]	AST[0]	FRE-AEt	FRE-PCA	PatchCore[0]	AST[0]	FRE
NVidia 2080-Ti	24s	7min 22s	5.9s	2.8s	18.4	19.6	313.1
Intel Xeon 8280	3mn 38s	25min 24s	1min 18s	6.3s	7.11	11.8	52.6
Intel i7-1270P	4mns 5s	5hr 8min	2min 52s	22.8s	1.28	2.8	16.4

Table 1: Training times and inference framerates evaluated on different computing platforms.

	features.1	features.2	features.3	features.4	features.5	features.6
Avg AUROC	74.7	86.6	89.5	95.0	97.6	98.6
Avg PRO	71.7	88.7	85.6	83.8	82.7	59.8
Avg pixel-wise AUROC	88.4	96.4	95.8	94.5	94.5	85.9

Table 2: FRE EfficientNet_B5 Performance across layers

	Layer1	Layer2	Layer3	Layer4
Avg AUROC	88.3	93.5	95.4	91.4
Avg PRO	85.3	90.5	90.2	68
Avg pixel-wise AUROC	94.3	96.8	97.3	92

Table 3: FRE Resnet18 Performance across layers

	Layer1	Layer2	Layer3	Layer4
Avg AUROC	87.9	90.7	96.0	91.9
Avg PRO	85.5	92.5	91.3	71.3
Avg pixel-wise AUROC	94.7	97.8	97.4	93.2

Table 4: FRE Resnet50 Performance across layers

	feat.8	feat.11	feat.13	feat.15	feat.18	feat.20	feat.22	feat.25	avgpool
Avg AUROC	84.0	88.8	89.5	90.2	90.1	90.8	92.2	93.0	91.6
Avg PRO	83.4	90.2	90.9	89.4	89.1	88.7	87.1	83.2	NA
Avg Pixel-wise AUROC	93.9	97.3	97.4	96.7	96.7	96.3	95.7	95.6	NA

Table 5: FRE VGG16 Performance across layers





Figure 1: Using a Resnet50 backbone, for each row: input image, ground truth segmentation mask, FRE segmentation map obtained with layer1, FRE segmentation map obtained with layer2, FRE segmentation map obtained with layer3, FRE segmentation map obtained with the combination of the 3 layers (layer1, layer2 and layer3).



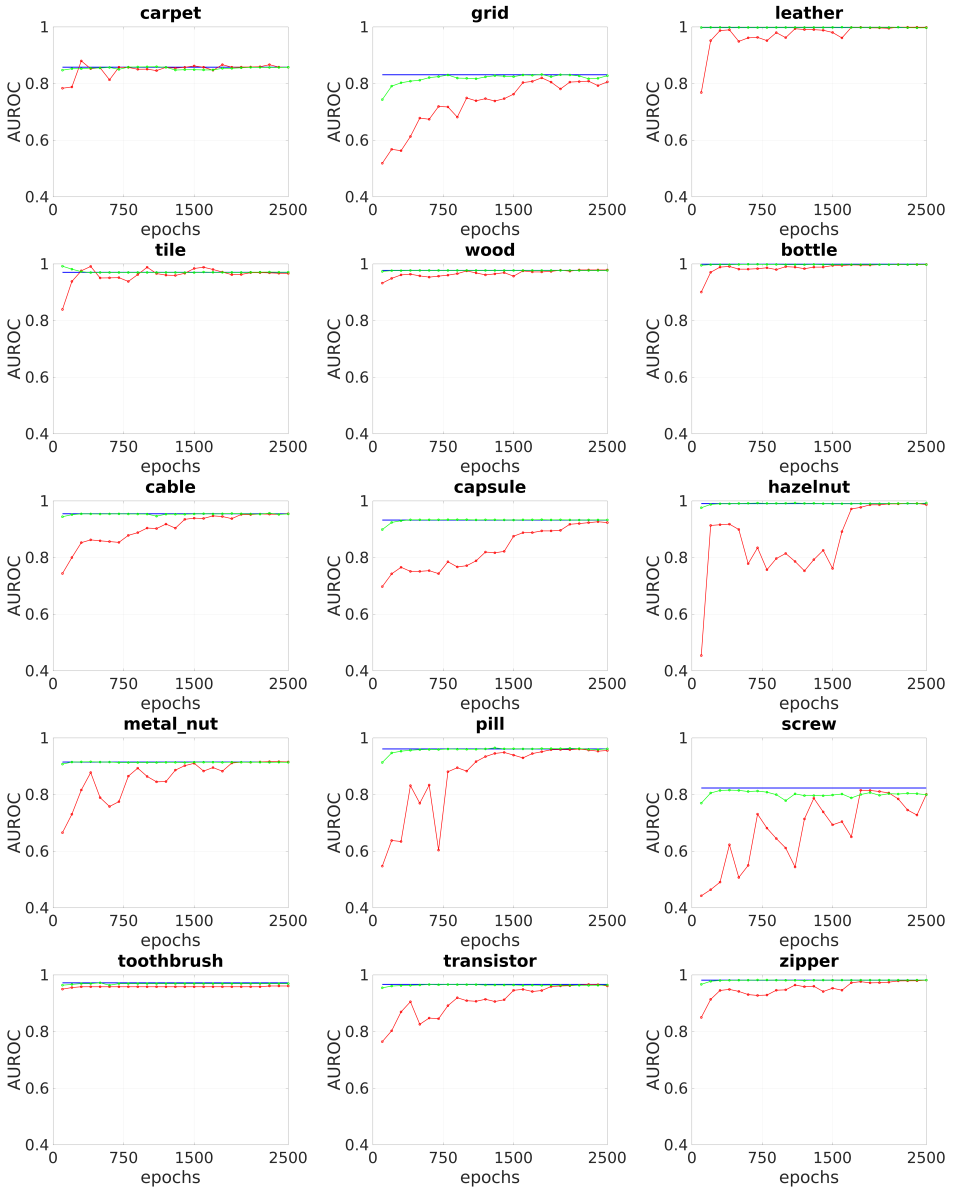


Figure 3: Class-itemized AUROC detection results for FRE-PCA (blue), FRE-AE (red) and FRE-AEt (green) with Resnet50 layer3. For FRE-AE and FRE-AEt, snapshots of the models being learned after every 100 epochs are used to show the convergence to FRE-PCA.

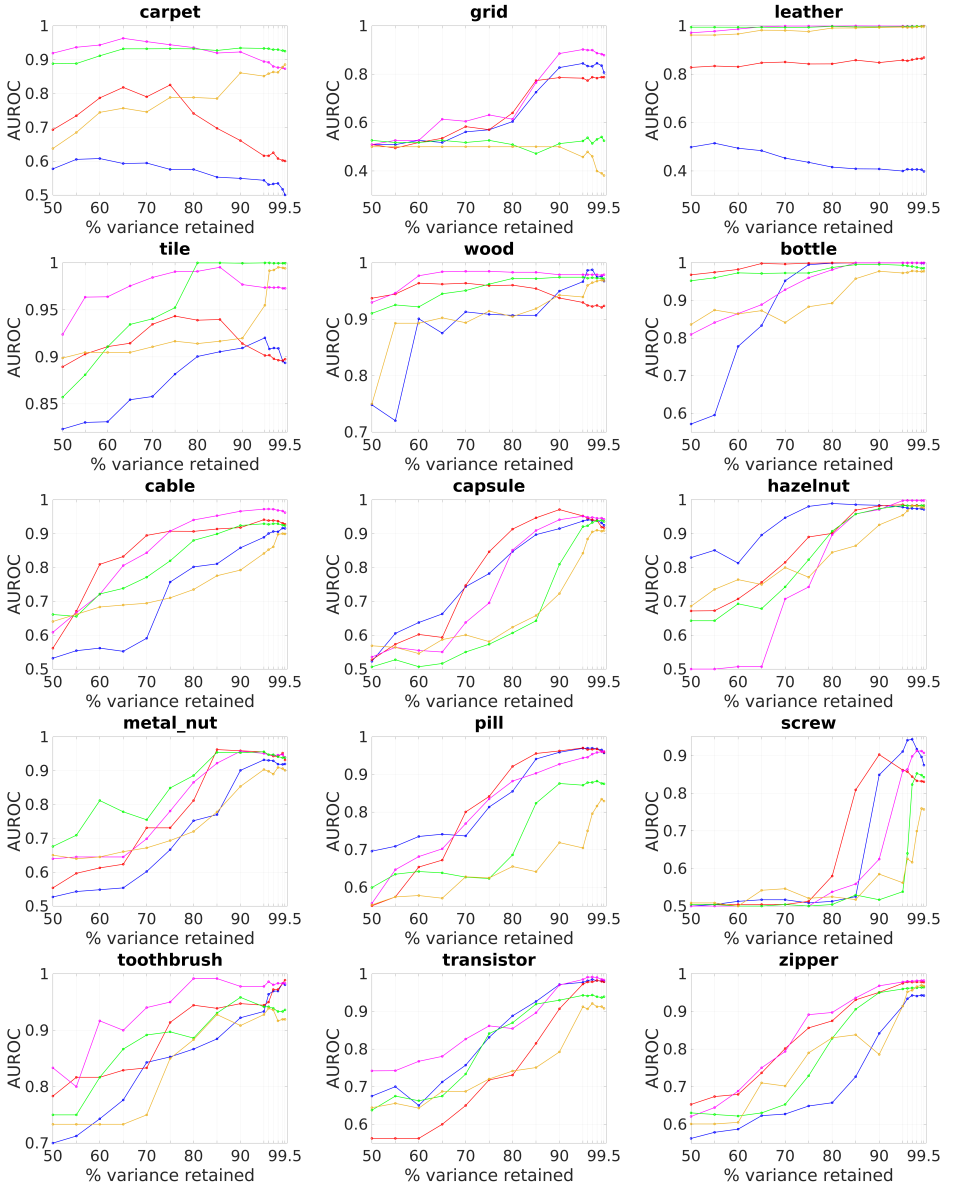


Figure 4: Class-itemized detection results (AUROC) while varying the only parameter of our method used with a ResNet50 backbone. Blue, red, magenta, green and gold resp. denote the following layers: layer1, layer2, layer3, layer4, and avgpool. They show insensitivity to parametrization, esp. within reasonable bounds [0.95, 1.0)

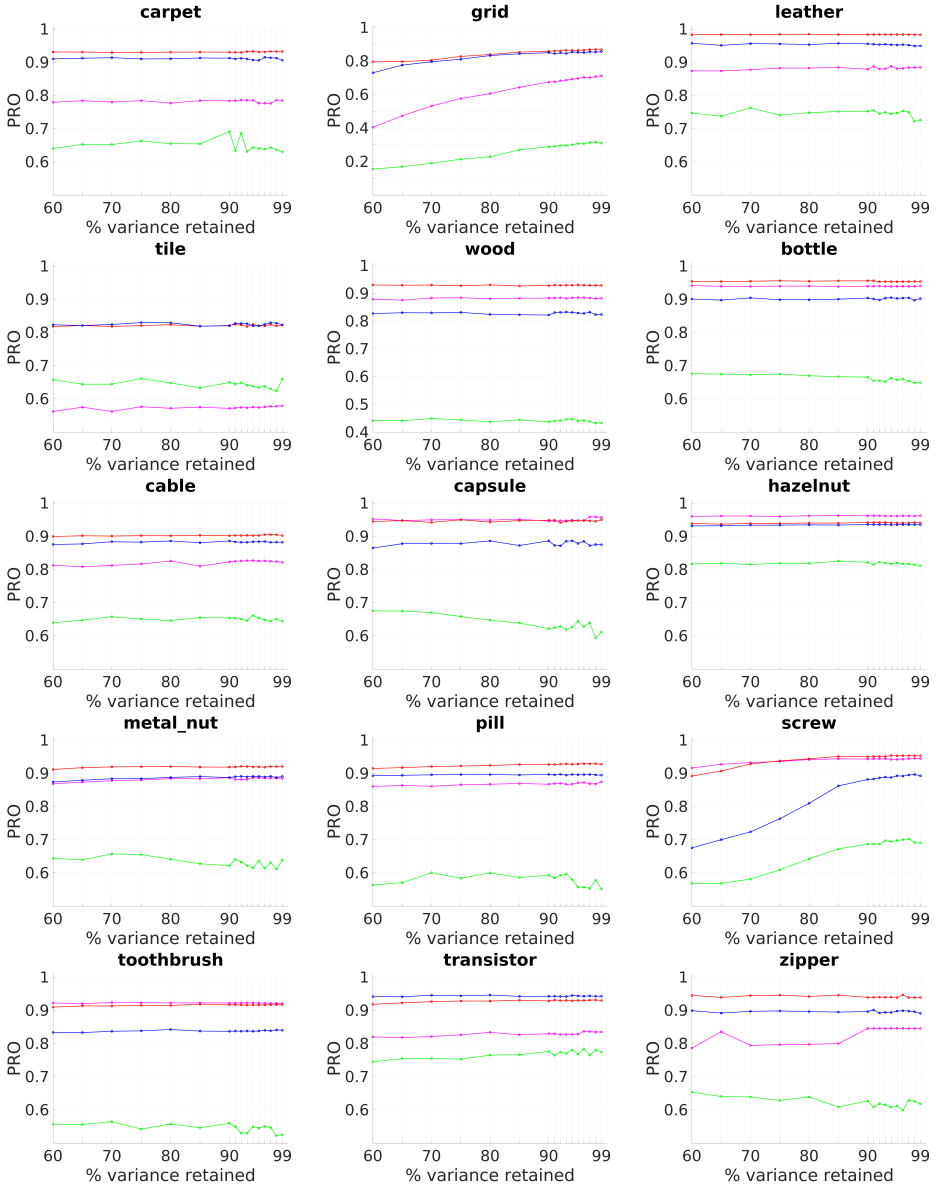


Figure 5: Class-itemized segmentation results (PRO) while varying the only parameter of our method used with a ResNet50 backbone. Blue, red, magenta, and green resp. denote the following layers: layer1, layer2, layer3, and layer4. They show insensitivity to parametrization, esp. within reasonable bounds [0.95, 1.0)

	CNN (Dict) [9]	AE (L2) [9]	SPADE [9]	PaDiM [9]	PatchCore [9]	AST [9]	FRE (1L & 3L)	
Carpet	72	59	97.5	99.1	99	-	98.3	99.2
Grid	59	90	93.7	97.3	98.7	-	97.3	98.1
Leather	87	75	97.6	99.2	99.3	-	99.8	99.8
Tile	93	51	87.4	94.1	95.6	-	95.1	96.5
Wood	91	73	88.5	94.9	95	-	96.7	97.7
Bottle	78	86	98.4	98.3	98.6	-	98.7	98.8
Cable	79	86	97.2	96.7	98.4	-	97.4	97
Capsule	84	88	99	98.5	98.8	-	98.8	99.1
Hazelnut	72	95	99.1	98.2	98.7	-	98.7	99.1
Metal nut	82	86	98.1	97.2	98.4	-	96.6	97.3
Pill	68	85	96.5	95.7	97.4	-	96.9	97
Screw	87	96	98.9	98.5	99.4	-	98.9	99.2
Toothbrush	77	93	97.9	98.8	98.7	-	98.4	98.7
Transistor	66	86	94.1	98.5	96.3	-	96.3	96.5
Zipper	76	77	96.5	98.5	98.8	-	98.6	98.8
Average	78	82	96	97.5	98.1	95.0	97.8	98.2

Table 6: MVTec Anomaly Segmentation benchmark with pixel-wise AUROC metric.

	Efficientnet B5		VGG16		Resnet18		Resnet50	
	1L	3L	1L	3L	1L	3L	1L	3L
carpet	94.6	97.6	97.5	97.8	97.9	99.0	98.4	99.2
grid	95.0	97.3	98.2	98.5	96.7	97.4	97.3	98.1
leather	98.6	99.3	98.8	99.0	99.7	99.8	99.8	99.8
tile	93.4	95.5	92.6	92.7	88.8	93.3	95.1	96.5
wood	91.7	93.8	96.2	96.3	94.6	96.6	96.7	97.7
bottle	98.6	98.3	97.9	97.6	98.4	98.6	98.7	98.8
cable	96.3	96.7	95.7	95.8	95.9	96.7	97.4	97.0
capsule	98.7	98.5	99.2	99.2	98.3	98.9	98.8	99.1
hazelnut	98.3	98.1	98.8	98.9	98.3	99.0	98.8	99.1
metal nut	96.2	95.5	96.3	96.4	95.4	97.0	96.6	97.3
pill	96.4	96.1	97.5	97.6	96.1	97.0	96.9	97.0
screw	98.4	98.1	99.5	99.6	99.1	99.3	98.9	99.2
toothbrush	98.3	97.9	98.5	98.5	98.2	98.9	98.4	98.7
transistor	95.4	96.7	95.6	95.6	95.6	96.4	96.3	96.5
zipper	96.9	97.7	98.3	98.0	98.4	98.8	98.6	98.8
Average	96.4	97.2	97.4	97.4	96.8	97.8	97.8	98.2

Table 7: MVTec Anomaly Segmentation: pixel-wise AUROC for FRE across backbones

	Layer1	Layer2	Layer3	Layer4
carpet	78.48	93.07	94.21	78.10
grid	71.36	87.21	89.29	48.97
leather	88.10	98.31	97.56	79.24
tile	57.87	82.49	82.71	70.89
wood	88.25	92.66	89.58	53.39
bottle	94.09	95.45	93.22	76.90
cable	82.02	90.21	85.04	70.86
capsule	95.90	94.34	91.75	68.69
hazelnut	96.25	94.06	93.10	84.86
metal nut	88.46	92.00	91.49	74.40
pill	87.34	93.01	90.41	72.29
screw	94.39	95.44	93.88	74.63
toothbrush	92.13	91.58	89.93	59.14
transistor	83.14	92.86	94.72	84.62
zipper	84.53	94.63	93.07	71.96
Average	85.49	92.49	91.33	71.26

Table 8: PRO across Resnet50 layers

	Layer1	Layer2	Layer3	Layer4
carpet	92.01	98.35	98.68	96.11
grid	90.66	97.28	97.57	81.92
leather	96.12	99.76	99.49	98.11
tile	83.15	95.09	95.22	94.17
wood	94.57	96.67	95.75	86.28
bottle	97.93	98.71	98.27	95.71
cable	94.72	97.37	95.56	93.11
capsule	98.62	98.81	98.37	95.62
hazelnut	98.87	98.75	98.35	96.83
metal nut	95.82	96.63	96.66	92.18
pill	94.24	96.86	95.64	93.57
screw	98.46	98.94	98.27	94.79
toothbrush	98.05	98.37	98.22	92.66
transistor	92.30	96.30	96.54	91.44
zipper	94.89	98.61	98.63	95.51
Average	94.69	97.77	97.41	93.2

Table 9: pixel-wise AUROC across Resnet50 layers

References

- [1] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9592–9600, 2019.
- [2] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *CoRR*, abs/2005.02357, 2020. URL <https://arxiv.org/abs/2005.02357>.
- [3] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. *CoRR*, abs/2011.08785, 2020. URL <https://arxiv.org/abs/2011.08785>.
- [4] P. Napoletano, F. Piccoli, and R. Schettini. Anomaly detection in nanofibrous materials by cnn-based self-similarity. *Sensors*, 18, 2018.
- [5] K. Roth, L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, and P. V. Gehler. Towards total recall in industrial anomaly detection. *CoRR*, abs/2106.08265, 2021.
- [6] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Asymmetric student-teacher networks for industrial anomaly detection. In *Winter Conference on Applications of Computer Vision (WACV)*, January 2023.