FRE: A Fast Method For Anomaly Detection And Segmentation

Ibrahima J. Ndiour, Nilesh A. Ahuja, Utku Genc and Omesh Tickoo
Intel Labs

BMVC 2023

Abstract

This work proposes a fast and principled method for unsupervised anomaly detection and segmentation. Our method operates under the assumption of having access solely to anomaly-free training data while aiming to identify anomalies of an arbitrary nature on test data.
We make our code available at: https://intellabs.github.io/dfm

Contributions

The principal contributions of the work include:

- a generalized approach that utilizes a shallow linear autoencoder as a principled out-of-distribution detection method operating in the feature space produced by a pre-trained DNN.
- a solid theoretical foundation for the method establishing the ***feature reconstruction error*** (FRE) as a principled measure of uncertainty.
- simultaneous solving of image-level anomaly detection and pixel-level anomaly segmentation.
- multiple implementation strategies addressing concerns related to memory, computational complexity, and dataset size.
- extensive experimentation showing state-of-the-art quality, as well as speed, robustness and insensitivity to parameterization.

Preliminaries: linear auto-associative networks

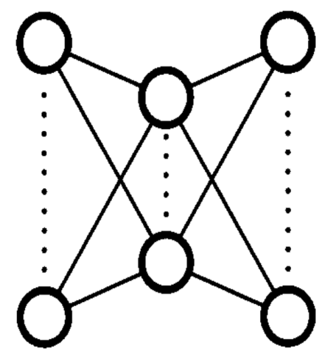

Auto-associative network.

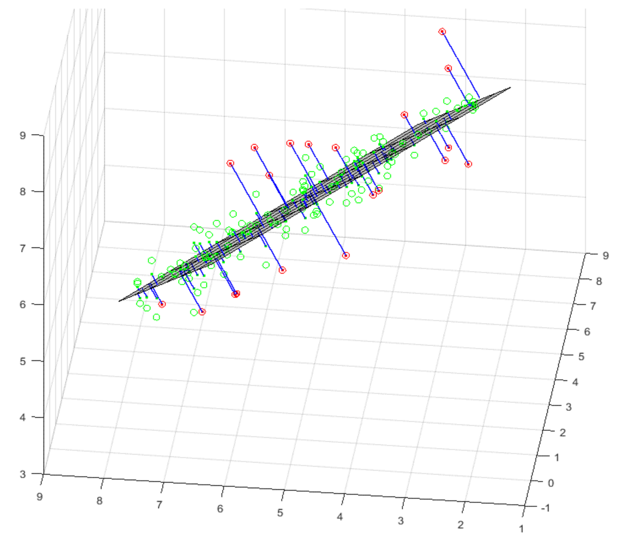

FRE visualization as distance to the principal subspace.

- The auto-associative network performs orthogonal projection onto the subspace spanned by the first principal eigenvectors of a covariance matrix associated with the training patterns.
- How does a neural network perform when exposed to a pattern never seen previously?
- In the auto-associative case, a precise quantitative answer can be given: the distortion on a new pattern is exactly given by its distance to the subspace generated by the first p eigenvectors of the data covariance matrix.

Proposed Approach

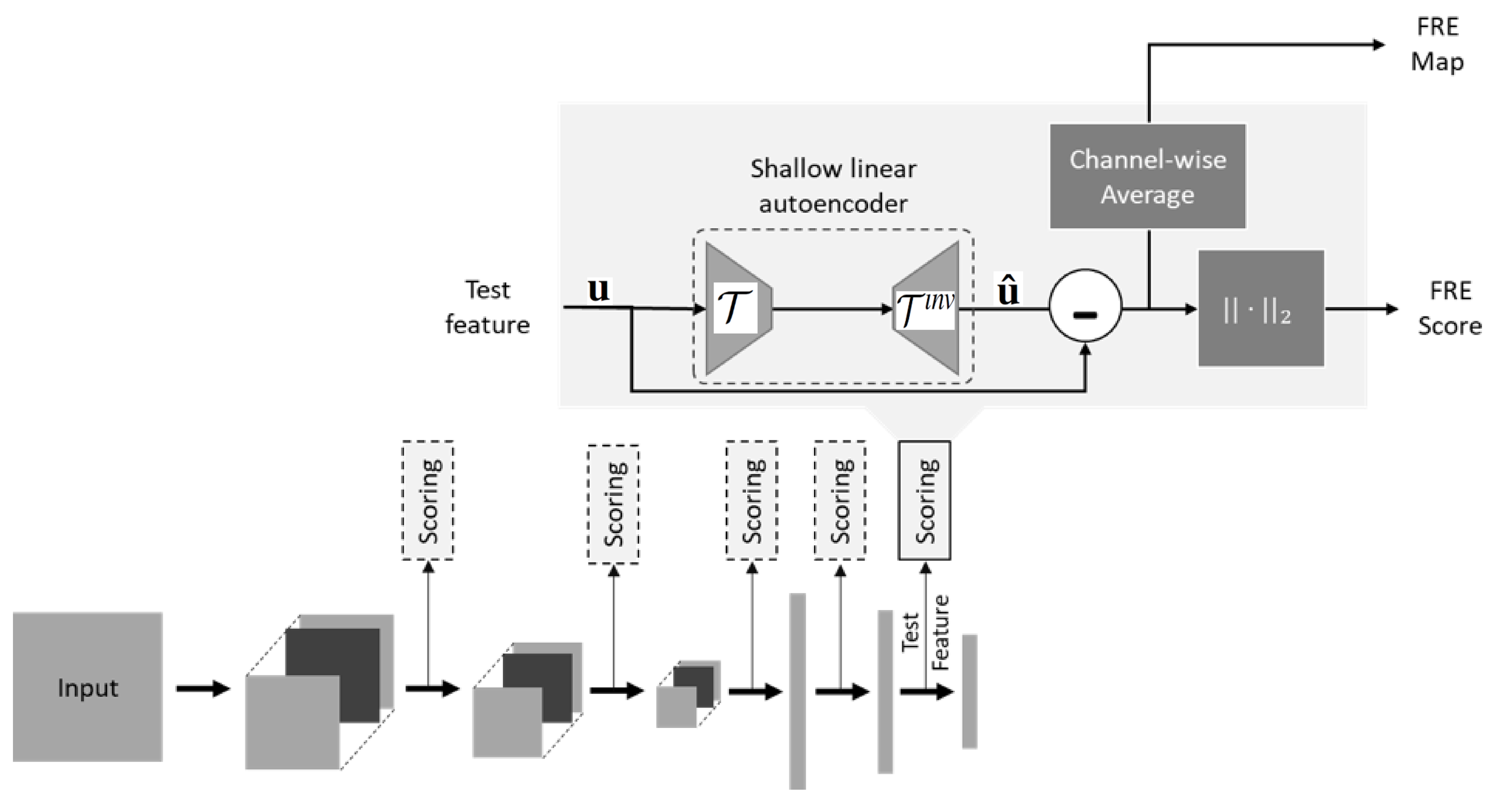

Block diagram of our proposed approach.

- **At training:** we pass the training data through the backbone network for a single forward pass and consider the intermediate (training) features at a given layer k. Subsequently, we train a (tied) shallow linear autoencoder with such features.
- **During inference:** the autoencoder is applied to a test feature $\mathbf{u}$ to obtain the reduced-dimension embedding, and its reconstruction back into the original (feature) space. The *feature reconstruction error* (FRE) is then calculated and used as an uncertainty score:

$$FRE(\mathbf{u},\mathcal{T}) = \mathbf{e} \triangleq \mathbf{u} - \hat{\mathbf{u}} = \mathbf{u} - (\mathcal{T}^{inv} \circ \mathcal{T})\mathbf{u}$$

- The following detection score is highly effective at discriminating between normal and anomalous samples:
- To derive the segmentation map, we perform channel-wise averaging on the FRE (re-arranged as a tensor) in order to accumulate the FRE errors along the channel dimension. This produces a single-channel FRE anomaly map M, where higher intensity regions correspond to anomalies.

$$\mathbf{M}_k(i,j) = \frac{1}{C_k}\sum_{c=1}^{C_k} \mathbf{e}(c,i,j)$$

Anomaly detection results

Category	GANomaly[1]	DifferNet[22]	SPADE* [10]	PaDiM*[11]	PatchCore[21]	AST[23]	FRE (Ours)
Average	76.2	94.9	85.5	97.9	99.1	**99.2**	98.6

MVTec dataset (AUROC metric ↑).

GANomaly[1]	1-NN[18]	DifferNet[22]	PatchCore[21]	FRE (Ours)
76.6	80.0	97.7	97.9	**99.2**

Magnetic Tile dataset (AUROC metric ↑).

ResNet18	ResNet50	VGG16	EfficientNet-B5
95.4	96.0	91.6	98.6

MVTec dataset for FRE across backbones (AUROC metric ↑).

Anomaly segmentation results

	CNN Dict	AutoEnc-L2	SPADE	PaDiM	PatchCore	AST	FRE (Ours)	
							1L	3L
Pixel AUROC (↑)	78	82	96	97.5	98.1	95	97.8	**98.2**
PRO (↑)	51.5	79	91.7	92.1	93.4	-	92.5	**93.9**

Pixel-AUROC and PRO metrics on MVTec dataset.

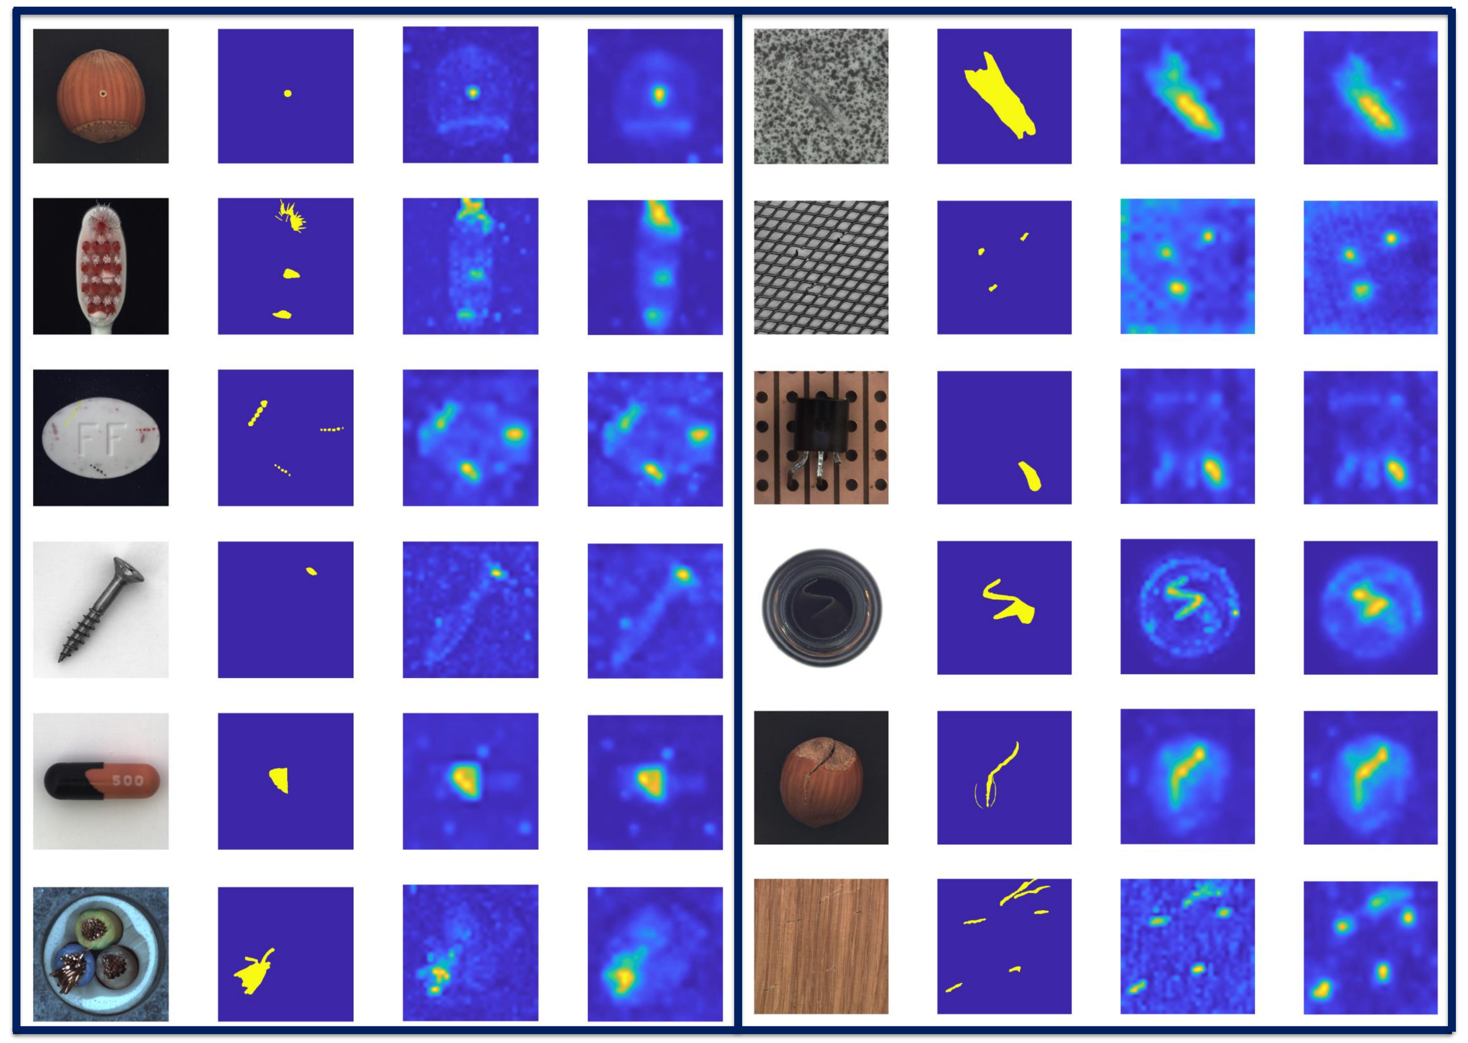

From left to right, each set of four images comprises: original image, ground truth segmentation mask, anomaly heatmap using FRE (our method) from a single layer, anomaly heatmap using FRE from three layers.

	Efficientnet B5		VGG16		Resnet18		Resnet50	
	1L	3L	1L	3L	1L	3L	1L	3L
Pixel AUROC (↑)	96.4	97.2	97.4	97.4	96.8	97.8	97.8	98.2
PRO (↑)	88.7	91.4	90.9	91.0	90.5	92.7	92.5	93.9

Pixel-wise AUROC and PRO metrics on MVTec for FRE across backbones.

Performance study

	Training Time (↓)				Inference frames-per-sec (↑)		
	PatchCore[5]	AST[6]	FRE-AEt	FRE-PCA	PatchCore[5]	AST[6]	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**

Training times and inference framerates evaluated on different computing platforms.

Conclusion

This work presented a fast, generalized approach for visual anomaly detection and segmentation. We propose applying a shallow, linear autoencoder on the intermediate features produced by a pretrained DNN and computing the *feature reconstruction error* (FRE) for use as uncertainty score. Our method meets or exceeds the state of the art in quality, is fast and requires no tedious manual tuning of parameters.