Anomaly Detection Using Convolutional Sparse Models

Diego Carrera, Giacomo Boracchi
Dipartimento di Elettronica, Informazione e Bioingegeria,
Politecnico di Milano, Italy
{giacomo.boracchi, diego.carrera}@polimi.it

Alessandro Foi
Department of Signal Processing
Tampere University of Technology, Finland
alessandro.foi@tut.fi

Brendt Wohlberg
Theoretical Division,
Los Alamos National Laboratory
Los Alamos, NM, USA
brendt@lanl.gov

We address the problem of detecting *anomalous* regions in images, i.e. regions having a structure that does not conform to *normal* images in a reference set [1]. Our approach is based on convolutional sparse models [2], which model an image $\mathbf{s} \in \mathbb{R}^{n_1 \times n_2}$ as the sum of k convolutions between filters $\mathbf{d}_m \in \mathbb{R}^{h_1 \times h_2}$ and sparse feature maps $\mathbf{x}_m \in \mathbb{R}^{n_1 \times n_2}, m \in \{1, \ldots, k\}$, i.e.

$$\mathbf{s} \approx \sum_{m=1}^{k} \mathbf{d}_m * \mathbf{x}_m \ . \tag{1}$$

Feature maps $\{\mathbf{x}_m\}$ of an input image $\mathbf{s}$ are computed by a sparse coding algorithm solving the optimization problem [2]

$$\underset{\{\mathbf{x}_m\}}{\arg\min} \frac{1}{2} \left\| \sum_m \mathbf{d}_m * \mathbf{x}_m - \mathbf{s} \right\|_2^2 + \lambda \sum_m \|\mathbf{x}_m\|_1 \,, \tag{2}$$

where $\lambda > 0$ is the parameter balancing the reconstruction error and the sparsity of the feature maps, and $\|\mathbf{d}_m\|_2 = 1 \, \forall m$. Solutions of (2) can be obtained via the Alternating Direction Method of Multipliers (ADMM) algorithm [3], exploiting an efficient formulation [4] in the Fourier domain. Filters can be learned from training images [2].

Our contribution is the design of an anomaly-detection approach monitoring also the local group sparsity of the feature maps, which is shown to be a relevant prior for anomaly detection.

I. Anomaly Detection

In our approach, filters $\{\mathbf{d}_m\}$ are learned to characterize the local structure of anomaly-free images represented by a training set S, as shown in Figure 1a. Then, in anomalous regions, where filters are less likely to match the image structures, it is reasonable to expect the sparse coding to be less successful, and that either the feature maps would be less sparse or that (1) would be a poor approximation.

Our intuition is that the sparsity measured as $\sum_m \|\mathbf{x}_m\|_1$ is too loose a criterion for discriminating anomalous regions, and that the distribution of nonzero coefficients across different feature maps should be also taken into account. In fact, we observed that within normal image regions, where filters are well matched with image structures, only a few feature maps are simultaneously active within a local subregion. In contrast, where filters and image structures do not match, more filters are typically active. For this reason, we perform anomaly detection by monitoring also the local group sparsity of the feature maps. Figure 1b illustrates feature maps within normal and anomalous regions.

To detect anomalous regions we analyze an input image $\mathbf{s}$ and the corresponding features maps $\{\mathbf{x}_m\}$ in a patch-wise manner. In particular, at first we compute the sparse coding of $\mathbf{s}$ yielding $\{\mathbf{x}_m\}$, then each patch of $\mathbf{s}$ is independently analyzed, computing an indicator to quantitatively assess the extent to which the patch is consistent with learned filters $\{\mathbf{d}_m\}$:

$$\mathbf{g}(i) = \begin{bmatrix} \left\| P_{i,q} \left(\mathbf{s} - \sum_m \mathbf{d}_m * \mathbf{x}_m \right) \right\|_2^2 \\ \sum_m \|P_{i,q} \mathbf{x}_m\|_1 \\ \sum_m \|P_{i,q} \mathbf{x}_m\|_2 \end{bmatrix} , \tag{3}$$

where $P_{i,q}$ denotes the matrix extracting a $q \times q$ patch centered at pixel i. The first and second elements of (3) represent the local reconstruction error and local sparsity of the feature maps respectively: these directly refer to the terms in (2), thus inherently indicate how successful the sparse coding was. The third element in (3) represents the group sparsity, and indicates the spread of nonzero coefficients across different feature maps in the vicinity of pixel i.

Indicators (3) are treated as random vectors and patches yielding outliers are considered anomalous. Outliers are simply detected as vectors falling outside a confidence region built around the mean vector $\overline{\mathbf{g}}$, previously computed from S. This yields the anomaly-detection criteria reported at line 6 of the following algorithm:

Training on a set of normal images S:
1. Learn filters $\{\mathbf{d}_m\}$, compute $\overline{\mathbf{g}}, \Sigma$ from $\mathbf{g}(i)$ extracted from S.
2. Set a threshold $\gamma > 0$.

Anomaly Detection in a test image $\mathbf{s}$:
3. Compute the filter maps $\{\mathbf{x}_m\}$ solving (2).
4. **foreach** pixel i of $\mathbf{s}$ **do**
5. Compute $\mathbf{g}(i)$
6. **if** $\left(\mathbf{g}(i) - \overline{\mathbf{g}}\right)' \Sigma^{-1} \left(\mathbf{g}(i) - \overline{\mathbf{g}}\right) > \gamma$ **then**
7. i belongs to an anomalous region
end
end

II. Experiments and Conclusions

We consider the following anomaly-detection approaches:

- **convolutional group**: anomaly detection monitoring $\mathbf{g}$ (3).
- **convolutional**: anomaly detection monitoring the first two components of $\mathbf{g}$ only.
- **patch-based**: each patch is independently approximated by a standard sparse model [5] rather than a convolutional one and both reconstruction error and sparsity are monitored, as in [6].

The first and second approaches use the same collections of filters (8 filters of size 8×8 and 8 of size 16×16) while dictionaries in the patch-based approach are 1.5 times redundant, and learned by [7].

We consider 25 textures from the Brodatz dataset [8]: the left half of each image is used for learning filters (or dictionary) and for computing $\overline{\mathbf{g}}$ and Σ. The right half of each image is used for preparing 600 test images as horizontal concatenation of two different textures (see the example in Figure 1b). Test images are processed using filters (or dictionary) learned from the corresponding left-hand side: thus any detection in the left half represents a false positive, and any detection in the right half a true positive. Performance is assessed from the receiver operating characteristic (ROC) curves in Figure 2, obtained varying γ. The regularization λ was empirically set to 0.1 and the patch size is $q \times q = 15 \times 15$.

Experiments indicate that the local group sparsity is an effective criterion for detecting anomalies in image structure. Ongoing work concerns the design of a specific sparse-coding algorithm including local group sparsity as a regularization term, and the use of this information for different imaging problems.

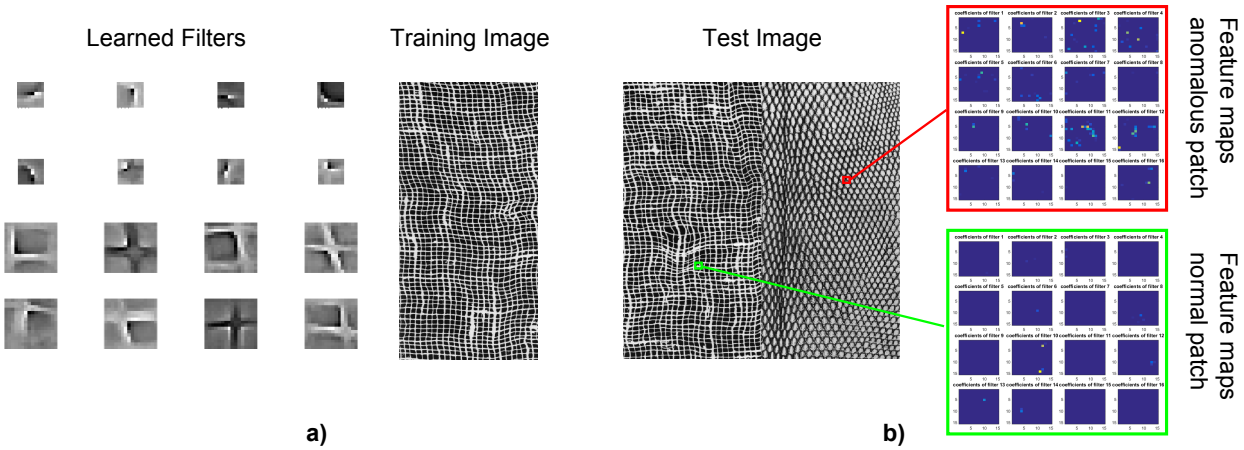

Fig. 1. (a) Learned filters (8 filters are of size 8×8 and 8 are of size 16×16) report the prominent local structures of the training image. (b) A test image used in our experiments: the left half represents the normal region (filters were learned from the other half of the same texture image), while the right half represents the anomalous region. The ideal anomaly detector should mark all pixels within the right half as anomalous, and all pixels within the left half as normal. The feature maps corresponding to the two highlighted regions (red and green squares) have approximately the same ℓ^1 norms. However, the maps at the right show that there is a substantially different spread of nonzero coefficients across feature maps. Here, the local group sparsity of the feature maps is more informative than sparsity for anomaly detection. Feature maps were rescaled for visualization sake.

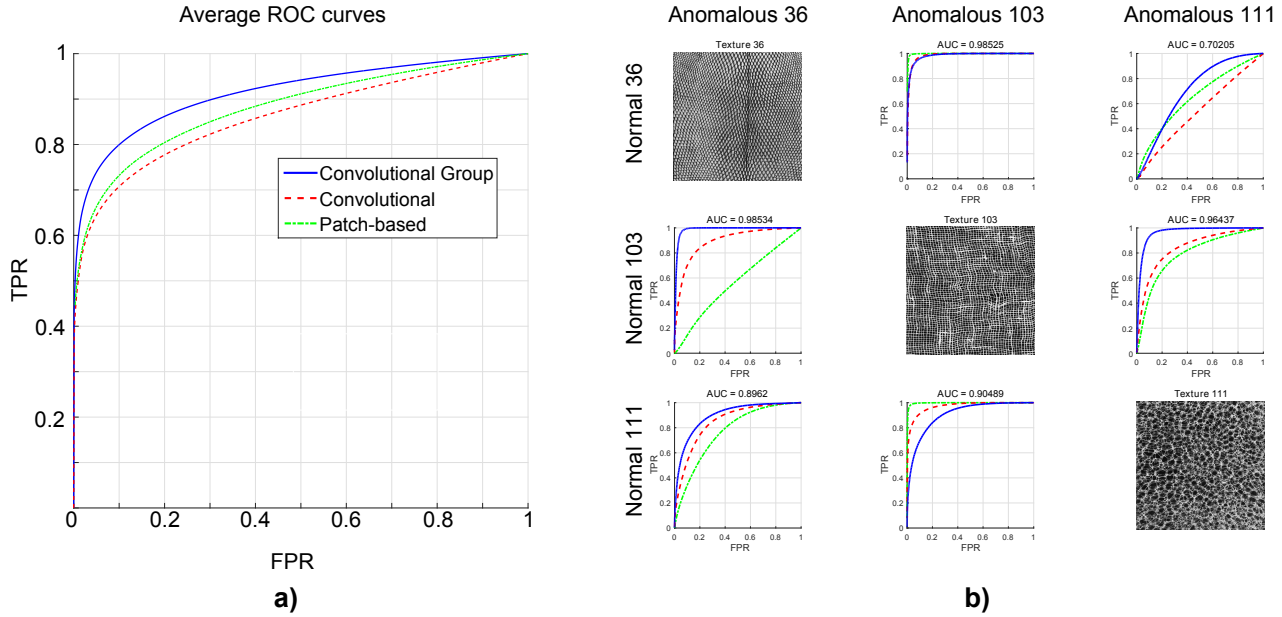

Fig. 2. (a) The average ROC curves of the approaches discussed in Section II over the 600 test images created from 25 textures from Brodatz dataset (textures 15, 20, 24, 27, 34, 36, 37, 49, 51, 52, 54, 55, 56, 65, 66, 68, 74, 76, 78, 83, 87, 103, 105, 109, 111). The Area Under the Curve values are 0.914, 0.861 and 0.881 for Convolutional Group, Convolutional and Patch-based, respectively. (b) Few ROC curves for selected textures. Anomaly detection performance using convolutional models (1) can be substantially improved monitoring also the group sparsity (blue solid line), as the gap between the blue and red curve (convolutional model) shows. The proposed approach outperforms also the patch-based sparsity (dotted green line). The subplots (b) show that while the anomaly-detection performance of the three solutions varies quite substantially among different images, the proposed approach performs satisfactorily in all these cases.

REFERENCES

[1] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, "A review of novelty detection," *Signal Processing*, vol. 99, pp. 215–249, 2014.

[2] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, "Deconvolutional networks," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 2528–2535.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[4] B. Wohlberg, "Efficient convolutional sparse coding," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 7173–7177.

[5] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM review*, vol. 51, no. 1, pp. 34–81, 2009.

[6] G. Boracchi, D. Carrera, and B. Wohlberg, "Novelty detection in images by sparse representations," in *Proceedings of IEEE Symposium on Intelligent Embedded Systems (IES)*, 2014, pp. 47–54.

[7] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *Proceedings of Annual International Conference on Machine Learning (ICML)*, 2009, pp. 689–696.

[8] P. Brodatz, *Textures: A Photographic Album for Artists and Designers*. Peter Smith Publisher, Incorporated, 1981.