
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Verdoja, Francesco; Grangetto, Marco
Graph Laplacian for Image Anomaly Detection

Published in:
Machine Vision and Applications

DOI:
[10.1007/s00138-020-01059-4](https://doi.org/10.1007/s00138-020-01059-4)

Published: 07/02/2020

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Verdoja, F., & Grangetto, M. (2020). Graph Laplacian for Image Anomaly Detection. *Machine Vision and Applications*, 31(1), Article 11. <https://doi.org/10.1007/s00138-020-01059-4>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Graph Laplacian for image anomaly detection

Francesco Verdoja¹ · Marco Grangetto²

Received: 9 October 2018 / Revised: 20 December 2019 / Accepted: 14 January 2020
© The Author(s) 2020

Abstract

Reed–Xiaoli detector (RXD) is recognized as the benchmark algorithm for image anomaly detection; however, it presents known limitations, namely the dependence over the image following a multivariate Gaussian model, the estimation and inversion of a high-dimensional covariance matrix, and the inability to effectively include spatial awareness in its evaluation. In this work, a novel graph-based solution to the image anomaly detection problem is proposed; leveraging the graph Fourier transform, we are able to overcome some of RXD's limitations while reducing computational cost at the same time. Tests over both hyperspectral and medical images, using both synthetic and real anomalies, prove the proposed technique is able to obtain significant gains over performance by other algorithms in the state of the art.

Keywords Anomaly detection · Graph Fourier transform · Graph-based image processing · Principal component analysis · Hyperspectral images · PET

1 Introduction

Anomaly detection is the task of spotting items that do not conform to the expected pattern of the data. In the case of images, it usually refers to the problem of spotting pixels showing a peculiar spectral signature when compared to all other pixels in an image. Image anomaly detection is considered one of the most interesting and crucial tasks for many high-level image- and video-based applications, *e.g.*, surveillance, environmental monitoring, and medical analysis [16].

One of the most used and widely validated techniques for anomaly detection is Reed–Xiaoli detector, often called RX detector for short [56], which is the most known example of covariance-based anomaly detectors. This class of detectors has found wide adoption in many domains, from hyperspectral [49] to medical images [65]; however, methods of this type suffer from crucial drawbacks, most noticeably the need for covariance estimation and inversion. Many situations exist where the drawbacks of these state-of-the-art anomaly

detectors lead to poor and unreliable results [67]. Moreover, the operations required by those techniques are computationally expensive [12]. For all these reasons, the research for a fast and reliable image anomaly detection strategy able to overcome the limitations of covariance-based anomaly detectors deserves further efforts.

In this paper, we use graphs to tackle image anomaly detection. Graphs are proved to be natural tools to represent data in many domains, *e.g.*, recommendation systems, social networks, or protein interaction systems [18]. Recently, they have found wide adoption also in computer vision and image processing communities, thanks to their ability to intuitively model relations between pixels. Graph-based approaches have been proposed to this date to solve a wide variety of image processing tasks, *e.g.*, edge detection [6], gradient estimation [55], and segmentation [9,59]. In particular, spectral graph theory has been recently bridged with signal processing, where the graph is used to model local relations between signal samples [57,60]. As an example, graph-based signal processing is emerging as a novel approach in the design of energy compacting image transformations [27,28,39,64,70].

To this date, graph-based approaches have not been proposed for image anomaly detection, although many techniques for anomaly detection on generic graphs have been explored in the literature [2]. Those techniques cannot be straightforwardly extended to images since they usually exploit anomalies in the topology of the graph to extract

✉ Francesco Verdoja
francesco.verdoja@aalto.fi

Marco Grangetto
marco.grangetto@unito.it

¹ School of Electrical Engineering, Aalto University, Maarintie 8, Espoo, Finland

² Department of Computer Science, University of Turin, Via Pessinetto 12, Turin, Italy

knowledge about the data [18]. On the other hand, in the image case the graph topology is constrained to the pixel grid, whereas different weights are assigned to edges connecting pixels depending on their similarity or correlation.

Our proposed approach uses an undirected weighted graph to model the expected behavior of the data and then computes the distance of each pixel in the image from the model. We propose to use a graph to model spectral or both spectral and spatial correlation. The main contribution of this paper is a novel anomaly detection approach which exploits spectral graph theory to overcome one of the well-known limitations of RX detector and other covariance-based anomaly detectors, *i.e.*, the need to estimate and invert a covariance matrix. Estimation of the covariance may be very critical in the presence of a small sample size; moreover, inverting such a matrix is also a complex, badly conditioned and unstable operation [40]. Our novel anomaly detector estimates the statistic of the background using a graph Laplacian matrix. Also, the graph model used by our approach is abstract and flexible enough to be tailored to any prior knowledge of the data possibly available. The effectiveness of our methodological contributions is shown in two use cases: a typical hyperspectral anomaly detection experiment and a novel application for tumor detection in 3D biomedical images.

The paper is organized as follows: We will first give a brief overview of RX detector and the graph Fourier transform in Sect. 2 and go over some related work in Sect. 3, and then we will present our technique in Sect. 4; we will then evaluate the performance of our technique and compare our results with those yielded by algorithms in the state of the art both visually and objectively in Sect. 5, and we will discuss these results in Sect. 6; finally, conclusions will be drawn in Sect. 7.

2 Background

Anomaly detection refers to a particular class of target detection problems, namely the ones where no prior information about the target is available. In this scenario, supervised approaches that try to find pixels which match reference spectral characteristics (*e.g.*, [24,42]) cannot usually be employed. This extends also to supervised deep learning or other data-driven approaches, which attempt to learn a parametric model from a set of labeled data. Although deep learning methods have found increasingly wide adoption for many other tasks in image processing and computer vision [15,35,71], their application to anomaly detection—especially on hyperspectral and medical imaging—is stifled by multiple factors: First, pixels have to be considered anomalous according to intra-image metrics which are difficult to capture in a dataset; second, the amount of data required to train the models is not often available in these contexts [11,44]. For these reasons, classical unsupervised

approaches are preferable instead. These algorithms detect anomalous or peculiar pixels showing high spectral distance from their surrounding [20]. To this end, the typical strategy is to extract knowledge of the background statistics from the data and then measure the deviation of each examined pixel from the learned knowledge according to some affinity function.

2.1 Reed–Xiaoli detector

The best known and most widely employed algorithm for anomaly detection is Reed–Xiaoli detector (RXD) by Reed and Yu [56]. To this date, it is still used as a benchmark algorithm for many anomaly detection applications [5,20,48,51]. RXD assumes the background to be characterized by a non-stationary multivariate Gaussian model, estimated by the image mean and covariance. Then, it measures the squared Mahalanobis distance [47] of each pixel from the estimated background model. Pixels showing distance values over a set threshold are assessed to be anomalous.

Formally, RXD works as follows. Consider an image $\mathbf{I} = [\mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_N]$ consisting of N pixels, where the column vector $\mathbf{x}_i = [x_{i1} x_{i2} \dots x_{im}]^T$ represents the value of the i th pixel over the m channels (or spectral bands) of \mathbf{I} . The expected behavior of background pixels can be captured by the mean vector $\hat{\boldsymbol{\mu}}$ and covariance matrix $\hat{\mathbf{C}}$ which are estimated as follows:

$$\hat{\boldsymbol{\mu}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i, \text{ and } \hat{\mathbf{C}} = \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{x}}_i \bar{\mathbf{x}}_i^T, \quad (1)$$

where $\bar{\mathbf{x}}_i = (\mathbf{x}_i - \hat{\boldsymbol{\mu}})$.

Mean vector and covariance matrix are computed under the assumption that vectors \mathbf{x}_i are observations of the same random process; it is usually possible to make this assumption as the anomaly is small enough to have a negligible impact on the estimate [12].

Then, the generalized likelihood of a pixel \mathbf{x} to be anomalous with respect to the model $\hat{\mathbf{C}}$ is expressed in terms of the square of the Mahalanobis distance [47], as follows:

$$\delta_{RXD}(\mathbf{x}) = \bar{\mathbf{x}}^T \hat{\mathbf{Q}} \bar{\mathbf{x}}, \quad (2)$$

where $\hat{\mathbf{Q}} = \hat{\mathbf{C}}^{-1}$, *i.e.*, the inverse of the covariance matrix, also known in the literature as the precision matrix.

Finally, a decision threshold η is usually employed to confirm or refuse the anomaly hypothesis. A common approach is to set η adaptively as a percentage of δ_{RXD} dynamic range as follows:

$$\eta = t \cdot p \cdot \max_{i=1, \dots, N} (\delta_{RXD}(\mathbf{x}_i)), \quad (3)$$

with $t \in [0, 1]$. Then, if $\delta_{RXD}(\mathbf{x}) \geq \eta$, the pixel \mathbf{x} is considered anomalous.

An interesting property of RXD has been observed by Chang and Heinz in [14]. In that work, the authors demonstrated how RXD can be considered an inverse operation of the principal component analysis (PCA).

More precisely, let us assume that $\kappa_1 \geq \kappa_2 \geq \dots \geq \kappa_m$ are the eigenvalues of the $m \times m$ covariance matrix $\hat{\mathbf{C}}$ and $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m\}$ is its set of unit eigenvectors with \mathbf{v}_j corresponding to κ_j . We can then form the matrix $\mathbf{V} = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_m]$ with the j th column specified by \mathbf{v}_j . \mathbf{V} can be used to decorrelate the signal by diagonalizing $\hat{\mathbf{C}}$ into the diagonal matrix \mathbf{K} whose j th diagonal element is κ_j , such that $\mathbf{V}^T \hat{\mathbf{C}} \mathbf{V} = \mathbf{K}$ and $\mathbf{V}^T \hat{\mathbf{Q}} \mathbf{V} = \mathbf{K}^{-1}$. Then, we can compute $\mathbf{y} = \mathbf{V}^T \bar{\mathbf{x}}$, which is known as the Karhunen–Loève transform (KLT). Data dimensionality reduction via PCA usually involves the computation of \mathbf{y} using just the first $p \ll m$ columns of \mathbf{V} . As shown in [14], (2) can be expressed as a function of \mathbf{y} as follows:

$$\begin{aligned} \delta_{RXD}(\mathbf{x}) &= \bar{\mathbf{x}}^T \hat{\mathbf{Q}} \bar{\mathbf{x}} \\ &= (\mathbf{V}\mathbf{y})^T \hat{\mathbf{Q}} (\mathbf{V}\mathbf{y}) \\ &= \mathbf{y}^T (\mathbf{V}^T \hat{\mathbf{Q}} \mathbf{V}) \mathbf{y} \\ &= \mathbf{y}^T \mathbf{K}^{-1} \mathbf{y} \\ &= \sum_{j=1}^m \kappa_j^{-1} y_j^2, \end{aligned} \quad (4)$$

where y_j represents the j th element of the KLT vector \mathbf{y} .

From this formulation, one can notice that RXD detects targets with small energies that are represented by small eigenvalues. This is because, according to (4), the smaller the eigenvalue, the greater its contribution to the value of δ_{RXD} . This is reasonable, since if an anomalous small target is present in the image, it will not be visible in the principal components, but it is rather going to appear in smaller components [12]. However, when seeing RXD in this form, it is quite evident that the last components, which are those containing mostly noise, are actually weighted the most. To improve the result of RXD, a value $p \ll m$ can be determined [38]. Then, the eigenvalues beyond the first (greater) p will be considered to represent components containing only noise and will be discarded. We then obtain a de-noised version of RXD that can be expressed as follows:

$$\delta_{RXD}^p(\mathbf{x}) = \sum_{j=1}^p \kappa_j^{-1} y_j^2. \quad (5)$$

Obviously, $\delta_{RXD}^m = \delta_{RXD}$.

The issue of determining p was addressed in [13,38] and is closely related to the problem of determining the intrinsic dimensionality (ID) of the image signal. Empiri-

cally, p is usually set such that a desired percentage $\psi \in [0, 1]$ of the original image cumulative energy content is retained. The cumulative energy content of the first p principal components of an image $\mathbf{I} = [\mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_N]$ can be expressed in terms of the image's KLT transform $\mathbf{Y} = \mathbf{V}^T \bar{\mathbf{I}} = [\mathbf{y}_1 \mathbf{y}_2 \dots \mathbf{y}_N]$ where $\bar{\mathbf{I}} = [\bar{\mathbf{x}}_1 \bar{\mathbf{x}}_2 \dots \bar{\mathbf{x}}_N]$ as

$$e(\mathbf{I}, p) = \sum_{i=1}^N \sum_{j=1}^p y_{ij}^2, \quad (6)$$

where y_{ij} is the j th element of the vector \mathbf{y}_i . We then choose the smallest $p \in [1, m]$, such that $e(\mathbf{I}, p)/e(\mathbf{I}, m) \leq \psi$. Commonly for dimensionality reduction applications $\psi = 0.9$, but for anomaly detection purposes that value might be too low, given we do not want to risk to lose the anomaly. In this case, $\psi = 0.99$ is usually more appropriate.

2.2 Graph Fourier transform

In recent years, the growing interest in graph-based signal processing [58] has stimulated the study of graph-based transform approaches. These methodologies map the image content onto a topological graph where nodes represent pixel intensities and edges model relations between nodes, *e.g.*, according to a criterion based on correlation or other similarity measures. The Fourier transform can be generalized to graphs obtaining the so-called graph Fourier transform (GFT) [57].

Consider an undirected, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ composed of a vertex set \mathcal{V} of order n and an edge set \mathcal{E} specified by (a, b, w_{ab}) , where $a, b \in \mathcal{V}$, and $w_{ab} \in \mathbb{R}^+$ is the edge weight between vertices a and b . Thus, a weighted graph can be described by its adjacency matrix \mathbf{W} where $\mathbf{W}(a, b) = w_{ab}$. A graph signal is a mapping that assigns a value to each vertex, denoted as $\mathbf{s} = [s_1 s_2 \dots s_n]^T$.

Typically, when computing the GFT a graph is constructed to capture the inter-pixel correlation and is used to compute the optimal decorrelating transform leveraging on spectral graph theory [60]. From the adjacency (also called weight) matrix \mathbf{W} , the *combinatorial graph Laplacian* matrix $\mathbf{L} = \mathbf{D} - \mathbf{W}$ can be computed, where \mathbf{D} is the degree matrix: a diagonal matrix whose a th diagonal element is equal to the sum of the weights of all edges incident to the node a . Formally,

$$\mathbf{D}(a, b) = \begin{cases} \sum_{k=1}^n w_{ak} & \text{if } a = b, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

In some scenarios, it is useful to normalize weights in the Laplacian matrix; in those cases, the use of the *symmetric normalized Laplacian matrix* \mathbf{L}^{sym} is preferred. It is defined

as

$$\mathbf{L}^{sym} = \mathbf{D}^{-\frac{1}{2}} \mathbf{L} \mathbf{D}^{-\frac{1}{2}}. \quad (8)$$

\mathbf{L}^{sym} has important properties, *i.e.*, its eigenvalues are always real, nonnegative, and bounded into the range $[0, 2]$; for these reasons, the spectrum of a symmetric normalized Laplacian relates well to other graph invariants for general graphs in a way that other definitions fail to do [18].

Any Laplacian matrix \mathbf{L} is a symmetric positive semi-definitive matrix with eigendecomposition:

$$\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T, \quad (9)$$

where \mathbf{U} is the matrix whose columns are the eigenvectors of \mathbf{L} and $\mathbf{\Lambda}$ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues. The matrix \mathbf{U} is used to compute the GFT of a signal \mathbf{s} as:

$$\tilde{\mathbf{s}} = \mathbf{U}^T \mathbf{s}. \quad (10)$$

The inverse GFT is then given by

$$\mathbf{s} = \mathbf{U} \tilde{\mathbf{s}}. \quad (11)$$

When computing the GFT, the eigenvalues in $\mathbf{\Lambda}$ are usually sorted for increasing magnitude, the first eigenvalue being equal to zero [57], *i.e.*, $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_m$. The eigenvectors in \mathbf{U} are sorted accordingly.

3 Related work

Despite its popularity, RXD has recognized drawbacks that undermine its performance in some applications. For a full discussion over the limitations of RXD, we suggest [12,67]; however, they can be summarized in the following:

1. RXD involves a high-dimensional covariance matrix that needs to be estimated and inverted, often under a small sample size [5,40]. Those are unstable, highly complex, and badly conditioned operations;
2. RXD often suffers from high false positive rate (FPR) [5,34,51];
3. RXD assumes that the background follows a multivariate Gaussian model, but there are cases in which this assumption might not be adequate, *e.g.*, in the case of multiple materials and textures [5,12,21,34];
4. RXD lacks spatial awareness: Every pixel is evaluated individually extrapolated from its context [31].

To address these issues, recent works have iterated over RXD's idea, *e.g.*, by considering subspace features [22,62],

by using kernels to go beyond the Gaussian assumption [21,41], by applying dimensionality reduction [33], by improving how the background statistics are estimated [20,50], or by exploiting sparsity and compress sensing theory [23,26,72]. In this work, we generalize RXD's idea by looking at it from the point of view of spectral graph theory. This not only makes us able to avoid costly covariance matrix inversions, but also allows us to incorporate spatial information and any prior knowledge about the background model into the detector. Previous work trying to including spatial awareness in the detector is available in the literature; a noteworthy example is whitening spatial correlation filtering (WSCF) [31], where the authors propose to apply a whitening transformation based on the eigendecomposition of the image covariance matrix. On the whitened space, RXD is represented by the Euclidean norm. Then, by using an approach based on constrained energy minimization, WSCF spots anomalous pixels by estimating consistency to their neighborhood in the whitened space. We compare our proposed approach to WSCF in the experimental section.

Although prior research targeting anomaly detection in graphs exists, it mostly focuses on anomalies in a graph structure, and not on graph signals [2,18]. For example, in the context of behavioral monitoring and intelligence, the structure of social graphs can be analyzed to spot subgraphs expressing patterns deviating from the rest of the network [52]. However, in images, the structure of the graph is fixed to a grid, and the application of graph-based anomaly detection algorithms coming from other domains is not straightforward; even in works where peculiarities in the graph signal are under observation, structure is included as part of the signal, as for example in [25] where a signal function of the physical distance between wireless sensors is proposed. The effectiveness of these approaches to images has not been reported yet.

Our proposed graph-based approach is founded on two recent findings: First, Zhang and Florêncio [70] have shown that a Laplacian model can be used as an estimation of the precision matrix \mathbf{Q} of an image, under the assumption that the image follows a Gaussian Markov random field (GMRF) model. This amounts to using a function of the partial correlation between nodes as graph weights. Second, it has been demonstrated how the GFT can be considered an approximation of the KLT for graph signals [39]. Recent literature in spectral graph theory has exploited this relationship to provide novel graph-based solutions to classical signal processing problems, in particular for image compression where the use of the GFT has been proposed as an alternative to the discrete cosine transform (DCT) [17,27,28,39]. This relationship is, however, never been explored in the context of image anomaly detection, which motivated us to study it in this work.

4 Method

In this work, we exploit the analogy between KLT and GFT in the framework of anomaly detection. In the GFT definition, the role of the covariance matrix in the KLT is taken by the graph Laplacian. It turns out that \mathbf{L} can be exploited also in the inverse problem of anomaly detection according to (4). We here propose a novel algorithm for image anomaly detection, which we will refer to as Laplacian anomaly detector (LAD). LAD overcomes some of the known limitations of RXD exposed in Sect. 2.1: It can be used to avoid problematic covariance matrix estimate and inversion, and it is able to include spatial information as well as a priori knowledge, when available.

4.1 Construction of the graph model

Given an image \mathbf{I} composed of N pixels and having m spectral bands or channels, we first build an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to serve as the model for the background pixels in the image. The graph is used to model local relations between pixel values and can be constructed to capture spectral and spatial characteristics. Topology and weights of the graph have to be chosen accordingly to the domain. We will discuss some general construction strategies in Sects. 4.3 and 4.4. The chosen graph will be described by a weight matrix \mathbf{W} , from which a Laplacian matrix \mathbf{L} will be computed according to the procedure detailed in Sect. 2.2. The use of the symmetric normalized Laplacian, constructed as in (8), in place of the unnormalized combinatorial one is to be preferred for the reasons expressed in Sect. 2.2. Also, \mathbf{L}^{sym} is proved to be preferable in similar domains, *e.g.*, segmentation and classification [7,29].

4.2 Graph-based anomaly detection

Given a pixel \mathbf{x} , we define a corresponding graph signal \mathbf{s} , *e.g.*, describing the spectral bands of \mathbf{x} or its spatial neighborhood, and compute the distance of \mathbf{x} from the model as

$$\begin{aligned} \delta_{LAD}(\mathbf{x}) &= \mathbf{s}^T \mathbf{L} \mathbf{s} \\ &= (\mathbf{U}\tilde{\mathbf{s}})^T \mathbf{L} (\mathbf{U}\tilde{\mathbf{s}}) \\ &= \tilde{\mathbf{s}}^T (\mathbf{U}^T \mathbf{L} \mathbf{U}) \tilde{\mathbf{s}} \\ &= \tilde{\mathbf{s}}^T \mathbf{\Lambda} \tilde{\mathbf{s}} \\ &= \sum_{j=1}^m \lambda_j \tilde{s}_j^2, \end{aligned} \quad (12)$$

where \tilde{s}_j represents the j th element of the GFT vector $\tilde{\mathbf{s}}$, and \mathbf{U} and $\mathbf{\Lambda}$ refer to the eigenvector and eigenvalue matrices used for the eigendecomposition of \mathbf{L} in (9). Although this formulation might look similar to the one of RXD given

in (4), some important differences have to be noted. First, the model used is not the inverse of the covariance matrix $\hat{\mathbf{C}}^{-1}$, but an arbitrary Laplacian model; this is a generalization over RXD, because if the image follows a GMRF model, then a Laplacian can be constructed to estimate the precision matrix [70], but if this is not the case a Laplacian model can be computed according to any knowledge of the domain. Second, the Laplacian matrix can be used to capture both spatial and spectral characteristics as we will detail in Sect. 4.4. Another thing to notice is that in (12) each contribution \tilde{s}_j is multiplied by λ_j , whereas in RXD each y_j was instead divided by the corresponding eigenvalue κ_j .

As already discussed for RXD, we can also use a de-noised version of the GFT where only the first smaller $p \ll m$ eigenvectors are kept, removing the higher and noisier frequencies and obtaining the following:

$$\delta_{LAD}^p(\mathbf{x}) = \sum_{j=1}^p \lambda_j \tilde{s}_j^2. \quad (13)$$

The parameter p is determined accordingly to the percentage of retained cumulative energy, following the approach presented in Sect. 2.1.

Finally, a decision threshold over δ_{LAD} is needed to determine if a pixel is anomalous or not. An approach similar to the one described in Sect. 2.1 can be employed.

4.3 Spectral graph model

As already mentioned, the graph model is used to characterize the typical behavior around the pixel being tested for anomaly. As in the case of standard RXD, the graph can be employed to model only the spectral relations: In this case, the vertex set \mathcal{V} consists of m nodes, each representing one of the spectral bands of \mathbf{I} ; then, we connect each pair of nodes (bands) with an edge, obtaining a fully connected graph. An example of this topology for a 3-band image is shown in

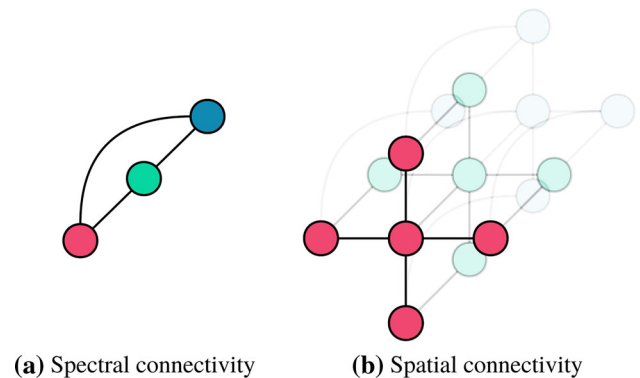


Fig. 1 Example of 3-band graph connectivity: the spectral components are fully connected, while spatially pixels are 4-connected

Fig. 1a. A weight is then assigned to each edge: If some a priori knowledge about inter-band correlation is available, it can be used to set weights accordingly; if this is not the case, a possibility is to use the image data to estimate the weights. Also, for each pixel \mathbf{x} , the graph signal \mathbf{s} will contain exactly the value of that pixel over the m bands, after removing the mean; thus, $\mathbf{s} = \bar{\mathbf{x}}$.

Under the assumption that the image follows a GMRF model, we might use partial correlation as weight, as proposed by Zhang and Florêncio [70]. To this end, given the precision matrix $\hat{\mathbf{Q}} = \hat{\mathbf{C}}^{-1}$, estimated according to (1), we can set the weight of the edge connecting nodes a and b as:

$$w_{ab} = -\frac{\hat{\mathbf{Q}}(a, b)}{\sqrt{\hat{\mathbf{Q}}(a, a) \hat{\mathbf{Q}}(b, b)}}. \quad (14)$$

Note that $w_{aa} = 0$ as we do not include self-loops. However, this approach still relies on the estimate and inversion of the covariance matrix that, as we already discussed, might be unreliable (especially in the presence of a small data sample) as well as expensive to compute: Matrix inversion requires $O(m^3)$ time [46]. Also, if the image does not follow a GMRF model, this distance function might produce unreliable results, as for all other covariance-based methods. An option to safeguard against this could be to use the graph constructed to evaluate the GMRF hypothesis with an approach similar to the one proposed in [3]

Another possibility is to use a different weight function, e.g., the Cauchy function [32], which has been proved to be able to capture graph distances effectively for image signals and is commonly used as graph weight in other applications like image segmentation and compression [8,28]. We propose to set the weight of the edge connecting bands a and b , according to the band mean vector $\hat{\boldsymbol{\mu}} = [\mu_1 \mu_2 \dots \mu_m]^T$ estimated as in (1), as

$$w_{ab} = \frac{1}{1 + \left(\frac{\mu_a - \mu_b}{\alpha}\right)^2}, \quad (15)$$

where α is a scaling parameter. In this study, we decided to set $\alpha = \frac{1}{m} \sum_{i=1}^m \mu_i$, to normalize all values according to the mean range of the bands. The advantages of this approach are twofold: It avoids using unreliable correlation estimates and does not require matrix inversion, thus reducing the computational cost significantly.

Although other approaches to estimate graph weights might be devised, in this study we will limit the analysis to these ones.

4.4 Integration of spatial information in the graph

One of the advantages of using a graph-based approach is the flexibility of the model. For example, by augmenting the

graph topology to include edges connecting each node to nodes describing the same band for the neighboring pixels, as shown in Fig. 1b, one is able to include spatial information in the model. We will refer to this spatially aware version of LAD as LAD-S.

When considering the case of 4-connected nodes, the resulting graph will be composed of $5m$ nodes; therefore, the weight matrix \mathbf{W} , as well as the corresponding Laplacian matrix \mathbf{L} , will be a $5m \times 5m$ matrix. We can construct the weight matrix as follows:

$$\mathbf{W}(a, b) = \begin{cases} w'_{ab} & \text{if nodes } a, b \text{ represent different} \\ & \text{bands of the same pixel,} \\ w''_{ab} & \text{if nodes } a, b \text{ belong to the same} \\ & \text{band of 4-connected pixels,} \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

where w'_{ab} and w''_{ab} are some spectral and spatial correlation measures, respectively.

Then, to compute the distance of a pixel \mathbf{x} from the model, a graph signal \mathbf{s} is constructed concatenating the vector corresponding to \mathbf{x} and its 4-connected neighbors; also in this case, the mean value $\hat{\boldsymbol{\mu}}$ is subtracted. It follows that the vector \mathbf{s} will have length $5m$.

The spectral weights w'_{ab} can be estimated as proposed in the previous section. The weights w''_{ab} can be used to enforce a spatial prior: As an example in the following experimental analysis, we will set uniform spatial weights $w''_{ab} = 1$.

5 Experiments

To objectively evaluate LAD's performance, we selected a couple of scenarios in which the use of RXD has been proposed. The first one is hyperspectral remote sensing, which is one of the most common use cases for anomaly detection where the use of RXD is widely validated [49]; the second one is the domain of 3D volumetric segmentation of tumoral masses on positron emission tomography (PET) images, where we successfully explored the use of RXD in the past [10,63,65]. In these scenarios, we compare the performance of the proposed technique with those produced by RXD and, in the hyperspectral domain, also with random-selection-based anomaly detector (RSAD) [20] and WSCF [31]. RSAD employs multiple random selections of pixels to estimate the background statistics and then marks a pixel as anomalous by merging the output of the different runs by a majority voting approach. WSCF applies a whitening transformation to the input based on the image covariance matrix and then incorporates spatial information in the anomaly measure. This latter algorithm is of partic-

ular interest for our evaluation, to compare its performance against our own spatially aware methodology.

5.1 Hyperspectral remote sensing

Hyperspectral images find wide adoption in remote sensing applications, where hyperspectral sensors are typically deployed on either aircraft or satellites. The data produced by these sensors are a three-dimensional array or “cube” of data with the width and length of the array corresponding to spatial dimensions and the spectrum of each point as the third dimension.

5.1.1 Dataset

The dataset used in this study is composed of three hyperspectral scenes collected by the 224-band AVIRIS sensor. As a common practice [12], we discarded the 20 water absorption bands, *i.e.*, bands (108–112, 154–167, 224). The first scene was collected over Salinas Valley, California, and is characterized by high spatial resolution (3.7-meter pixels). The area covered by this scene comprises 512 lines by 217 samples, and it includes vegetables, bare soils, and vineyard fields. A classification ground truth containing 16 classes is provided with this scene. A sample band of the image together with the classification ground truth is shown in Fig. 2. The other two scenes image two urban environments and come with anomaly detection ground truth, and both comprise 100 lines by 100 samples. We will refer to them as *Urban-A* and *Urban-B*. A sample band of these two scenes, together with their corresponding ground truth, is shown in Fig. 3.




Fig. 2 The full 512×217 Salinas scene

To evaluate LAD, we tested it on both real and synthetic anomalies. For the Salinas scene, we cropped a 200×150 portion of the scene and manually segmented a construction which was visible in the cropped area: As the scene mostly contains fields of various kinds, this human-made construction was a good anomalous candidate. This setup, which we will call *Field*, is shown in Fig. 3m together with its ground truth in Fig. 3n.

To obtain a synthetic anomaly, we used the target implant method [61] on a different portion of the Salinas scene. The 150×126 binary mask image \mathbf{M} shown in Fig. 3t has been constructed by generating six squares having sides measuring from 1 to 6 pixels arranged in a line. The six squares have been then copied in reverse order and arranged in another line at close distance. The two lines have finally been rotated by an angle of approximately $\pi/6$. The pixels inside the squares have a value of 1, while the rest of the pixels in \mathbf{M} have value 0. Then, we cropped a region \mathbf{I} from the Salinas scene, having the same dimension as the mask. We used it to build the modified image \mathbf{I}' containing the implanted target as follows:

$$\mathbf{I}'(i, j) = \mathbf{M}(i, j) \cdot \Phi(k) + (1 - \mathbf{M}(i, j)) \cdot \mathbf{I}(i, j), \quad (17)$$

where Φ is a function that, given a parameter $k \in [1, 16]$, returns a random pixel from the region of the Salinas scene having class k according to the classification ground truth shown in Fig. 2b. In the following discussion, for conciseness, we will limit the analysis to two synthetic setups with $k = 14$ and $k = 4$, respectively. The two representative values have been chosen since RXD achieves the best performance on the former and the worst one on the latter. We will refer to them as *Impl-14* and *Impl-4*, respectively. A sample band from the *Impl-14* setup is shown in Fig. 3s.

Figure 4 shows the mean and standard deviation of the intensity of each band for the background, the anomaly region in *Impl-4* and *Impl-14*. As it can be noticed, the spectral characteristics of the anomaly in *Impl-4* are similar in shape to those of the background, although with reduced intensities. The anomaly in *Impl-14* presents a more different curve than the others, instead.

5.1.2 Experimental results

We are interested in evaluating the detection accuracy of LAD using the Laplacian model built over the partial correlation weights (\mathbf{L}_Q) and the one built using Cauchy distance (\mathbf{L}_C). Also, we want to test both the spectral version of LAD and its spatially aware variant LAD-S. The results will be compared with those yielded by classic RXD, RSAD, and WSCF. We compare our results against those yielded by RXD, given its well-known status as benchmark algorithm for anomaly detection. We want also to confirm with our




Fig. 3 Hyperspectral test scenarios and algorithm outputs. LAD results have been obtained using L_C




Fig. 4 Spectral characteristic curves for different regions of the image. The line represents the mean intensity computed over all pixels in a region, while the shaded area represents the standard deviation

experiments one of the known limitations of RXD enunciated in Sect. 2.1, namely how the inclusion of spatial information in RXD is detrimental to its performance, to demonstrate how our approach overcomes this limitation. Another well-known algorithm which aims at addressing this limitation is

WSCF, and for this reason we selected it for evaluation as well. WSCF requires a parameter α to determine the amount of spatial information included in the metric. In this study, we set $\alpha = 0.2$, as suggested in the original work [31]. RSAD requires to select: the initial number of randomly selected blocks N , which should be as small as possible but still large enough so that $4N > b$, where b is the number of image bands; the number of random selections L ; and the percentile α . For these parameters, we chose the following values in our experiments: $N = 80$, $L = 40$, and $\alpha = 0.001$. We implemented our method as well as all three benchmark methods in MATLAB 2014b. All experiments were run on a laptop equipped with an Intel® Core™ i7@2.20GHz CPU, a NVIDIA GT435M GeForce GPU and 8GB of RAM ¹.

Figure 3 shows the visual results by LAD (L_C) approach compared to the ones yielded by RXD, RSAD, and WSCF on

¹ The hyperspectral datasets and all algorithm implementations used for the experiments presented in this work can be found at: github.com/fverdoja/LAD-Laplacian-Anomaly-Detector.

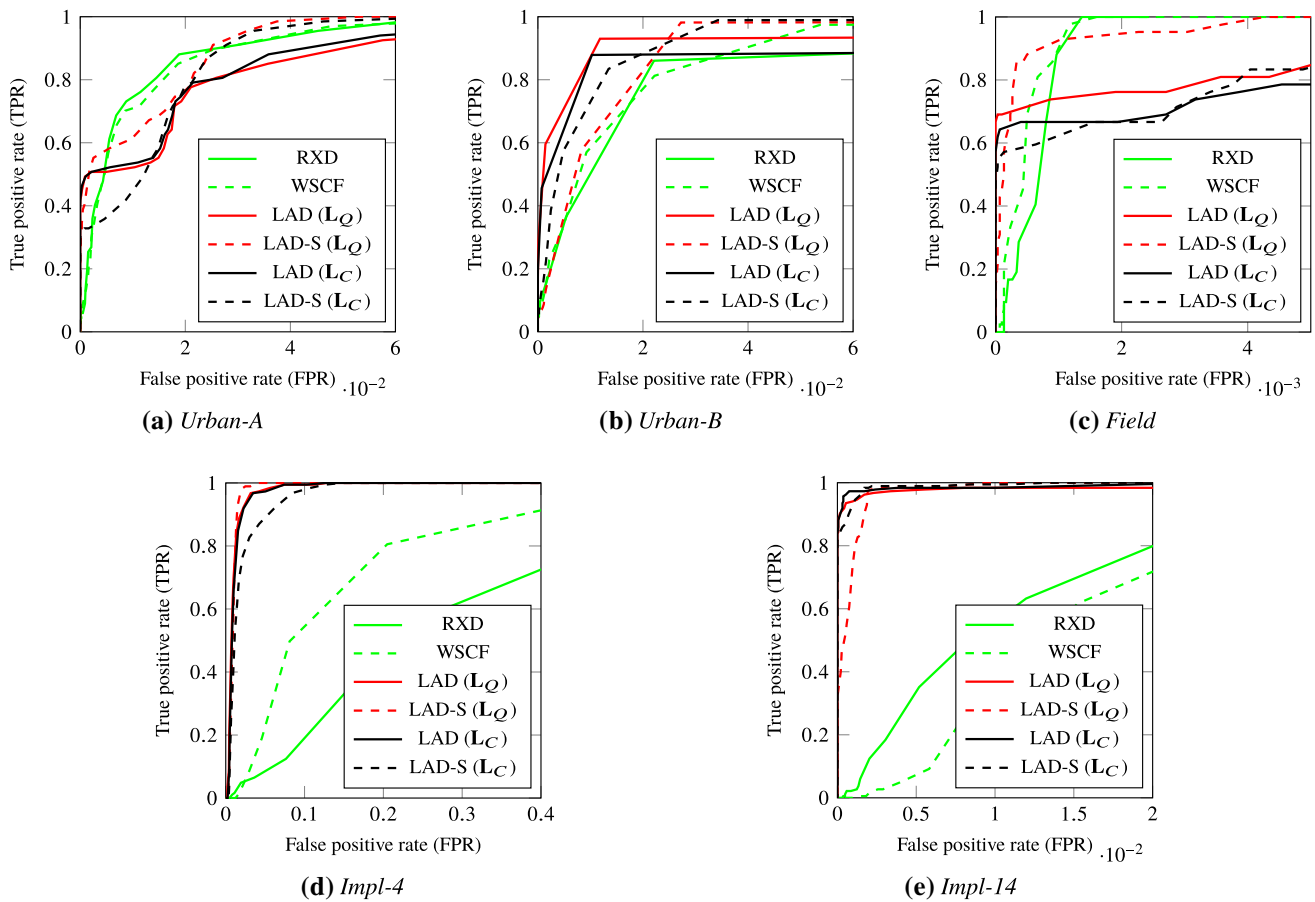


Fig. 5 Receiver operating characteristic (ROC) curves for the hyperspectral testing scenarios

the all hyperspectral scenarios. It can be clearly noticed that the lower number of false positives LAD is able to achieve against all other algorithms.

Figure 5 shows the ROC curves for the hyperspectral test cases, for all algorithms except RSAD. The approach by virtue of which RSAD selects which pixels are anomalous does not lend itself to be plotted in a ROC curve. The scale of the FPR axis has been enhanced, as common in anomaly detection studies [4,43,68], given the great difference in scale between the number of negative pixels and positive ones. It can be noticed how in all scenarios except *Urban-A* our approach outperforms both RXD and WSCF. On *Urban-A*, all algorithms perform very similarly. Also, worth noticing is that the inclusion of spatial information yields limited improvements on the hyperspectral scenarios. When comparing results obtained by LAD using \mathbf{L}_Q or \mathbf{L}_C , it can be noticed how performance is often very similar. This is a remarkable result, also considering that \mathbf{L}_C creates a model of the background without the need for matrix inversions, so it proves to be both quicker and equally precise.

To further compare performance yielded by the different approaches, we also use the standard spatial overlap index (SOI)

[73], also known as Dice similarity coefficient (DSC) [19], which can be computed as follows:

$$SOI = \frac{2(A \cap B)}{A + B}, \quad (18)$$

where A and B are two binary masks (*i.e.*, the ground truth or region of interest (ROI) and the output of an automatic algorithm); the intersection operator is used to indicate the number of pixels/voxels having value 1 in both masks, while the sum operator indicates the total number of pixels/voxels having value 1 in the two masks. SOI is also equivalent to the statistical F_1 -score, which is the harmonic mean of precision and sensitivity, and is usually defined in terms of Type I and Type II errors as follows:

$$F_1 = \frac{2 \cdot \text{true positive}}{2 \cdot \text{true positive} + \text{false positive} + \text{false negative}}. \quad (19)$$

The equality between (18) and (19) can be easily demonstrated considering that $A \cap B$ contains the true positive pixels/voxels and that if we consider that $A = (\text{true positive} +$

Table 1 Experimental results in hyperspectral setup (SOI)


	Urban-A	Urban-B	Field	Impl-14	Impl-4	Average
RXD	0.508	0.649	0.685	0.445	0.045	0.466
RSAD	0.078	0.310	0.042	0.450	0.022	0.180
WSCF	0.489	0.623	0.708	0.391	0.103	0.463
LAD (\mathbf{L}_Q)	0.606	0.791	0.806	0.941	0.525	0.734
LAD-S (\mathbf{L}_Q)	0.576	0.664	0.818	0.898	0.540	0.699
LAD (\mathbf{L}_C)	0.614	0.782	0.754	0.954	0.514	0.724
LAD-S (\mathbf{L}_C)	0.467	0.721	0.697	0.919	0.409	0.643

Bold values indicates the best result

Table 2 Experimental results after dimensionality reduction in hyperspectral setup (SOI)

	Urban-A	Urban-B	Field	Impl-14	Impl-4	Average	Gain (%)
RXD ^p	0.692	0.304	0.930	0.965	0.355	0.649	+39.19
LAD ^p (\mathbf{L}_Q)	0.606	0.791	0.806	0.941	0.521	0.733	−0.11
LAD-S ^p (\mathbf{L}_Q)	0.603	0.659	0.817	0.928	0.579	0.717	+2.57
LAD ^p (\mathbf{L}_C)	0.606	0.776	0.789	0.951	0.535	0.731	+1.08
LAD-S ^p (\mathbf{L}_C)	0.462	0.725	0.706	0.945	0.423	0.652	+1.49

Bold values indicates the best result

**Fig. 6** Energy and eigenvalue curves for the *Impl-14* scenario

false positive) and $B = (\text{true positive} + \text{false negative})$, then also the denominator in (18) equals the one in (19). Clearly, to compute the SOI metric one needs to select a threshold t to identify the anomaly subset B . Many approaches [1,53,69] have been proposed in the literature to deal with the problem of choosing the optimal threshold. In this work, we select the value of t yielding the highest SOI, *i.e.*, striking the best balance between TPR and FPR on the ROC curve in terms of SOI. This choice allows us to compute a single-objective metric to compare the analyzed methods. Alternatively, we could also use the area under the curve (AUC), which measures the area under each ROC curve; we decided to avoid such metric since it has been recently criticized for being sensitive to noise [36] and for other significant problems it shows in model comparison [37,45].

Table 1 shows all SOI results of our tests. It can be noticed how all variants of our approach are able to outperform RXD,

RSAD, and WSCF. These results are consistent with those presented by the ROC curves.

Finally, in Table 2 we show results of the de-noised version of both LAD and RXD, which we call LAD^p and RXD^p, respectively. In this case, the value of p has been chosen according to the cumulative energy as described in Sect. 2.1, setting $\psi = 0.99$. It can be noticed how RXD is able to gain the most from dimensionality reduction. These results can be explained considering the distribution of energy in the eigenspace decomposition. For the *Impl-14* scenario, in Fig. 6 we show the cumulative energy distribution in the different eigenspaces together with the corresponding eigenvalues κ_j^{-1} and λ_j (that are used to weigh the different contribution in (5) and (13), respectively). It can be noticed that in the RXD case (Fig. 6a) energy is better compacted into few eigenspaces with respect to LAD (Fig. 6b, c). At the same time, it can be observed that the distribution of κ_j^{-1} in RXD dramatically




Fig. 7 The three FDG-PET images of one of the sample patients; (1) is the early scan (ES, $144 \times 144 \times 213$ px), (2) and (3) are constructed integrating the delayed scan in 3-min time windows (DS1 and DS2, $144 \times 144 \times 45$ px). Only the area containing the tumor is acquired in the delayed scan. These images, originally in grayscale, are here displayed using a *Fire* lookup table

amplifies the last eigenspaces, *i.e.*, the noise components, according to (5). On the contrary, this phenomenon does not affect LAD since the distribution of eigenvalues λ_j is not peaked on the last eigenspaces. It follows that the effect of noise in (13) is mitigated by construction and the benefit of dimensionality reduction is limited. Indeed, it can be noted that results obtained by RXD after dimensionality reduction are in line with those obtained by LAD in its simple form. Being the eigendecomposition a costly operation, on a par with matrix inversion, the use of LAD (\mathbf{L}_C), which does not require any matrix inversion or eigendecomposition, might be preferable.

5.2 Application to 3D volumes: tumor segmentation in PET sequences

PET data are volumetric medical images that are usually employed to locate the tumoral area for proper oncological treatment, *e.g.* by means of radiotherapy. From a PET scan, one or more 3D images can be produced where the intensity of a voxel represents the local concentration of the tracer during the time window of the scan. In particular, fluorodeoxyglucose positron emission tomography (FDG-PET) is used to detect tissue metabolic activity by virtue of the glucose uptake.

During normal cell replication, mutations in the DNA can occur and lead to the birth of cancer cells. By their nature, these cells lack the ability to stop their multiplication, raising cell density in their region and causing insufficient blood supply. The resulting deficiency in oxygen (hypoxia) forces these cells to rely mostly on their anaerobic metabolism, *i.e.*, glycolysis [54]. For this reason, glycolysis is an excellent marker for detecting cancer cells; FDG-PET—in which the tracer's concentration indicates the glucose uptake in the imaged area—turns out to be a suitable tool for recognizing

tumors, metastases, and lymph nodes all at once [30]. It follows that proper segmentation of tumors in medical images is crucial as oncological treatment plans rely on precise information on the tumoral region to be effective [54]. Manual segmentation by medical staff has been proven to be subjective, inaccurate, and time-consuming [66]; for this reason, the need for automatic methods for tumor region segmentation is on the rise. PET images carry information about cells metabolism and are therefore suitable for this task; however, PET segmentation is still an open problem mainly because of limited image resolution and strong presence of acquisition noise [69].

In [10,63,65], we successfully explored the use of RXD to identify the anomalous behavior of cancer cells over time in sequences of three FDG-PET images acquired over a time span of one hour. A quick visual overview of this setup is shown in Fig. 7. The idea behind the use of RXD in this scenario arises from the fact that cancer cells tend to acquire glucose differently than normal cells, given their peculiar reliance on anaerobic metabolism. For this reason, when considering the values a voxel assumes over time, cancer's anomalous glucose uptake can be successfully spotted using anomaly detection techniques, where the usual role of spectral bands is taken by three PET images acquired over time.

To do this, we build a 4D matrix \mathbf{I} , having the three spatial dimensions as the first three dimensions and the time as the fourth dimension. Being acquired at different times, with the subject assuming slightly different positions, it is worth recalling that the images need to be aligned using registration algorithms as detailed in [65]. The resulting matrix \mathbf{I} will then have size $144 \times 144 \times 45 \times 3$. Then, for a generic voxel, identified by its spatial coordinates, we define the vector $\mathbf{x} = [x_1 x_2 x_3]^T$ as the vector containing that voxel's intensities over time. In other words, RXD can be employed in this case if time takes the role of the spectral dimension.

5.2.1 Experimental results

In this study, we used a dataset comprising eight patients, that has been made available by the Candiolo Cancer Institute (IRCCS-FPO) for research purposes. All the acquisitions have been made using a *Philips Gemini TF PET/CT*. To this end, we acknowledge the precious aid of nuclear medicine physicians who have manually segmented the ROIs on the PET images, setting up the ground truth for evaluating the performance yielded by the proposed tools. We will refer to this setup as *Tumor*.

Also in this scenario, we are interested in evaluating the detection accuracy of LAD using both Laplacian models, \mathbf{L}_Q and \mathbf{L}_C , and compare our results with those yielded by classic RXD. We cannot compare with WSCF in this domain as its extension to 3D has not been proposed, and therefore the choice of the parameter α is non-trivial. A thing

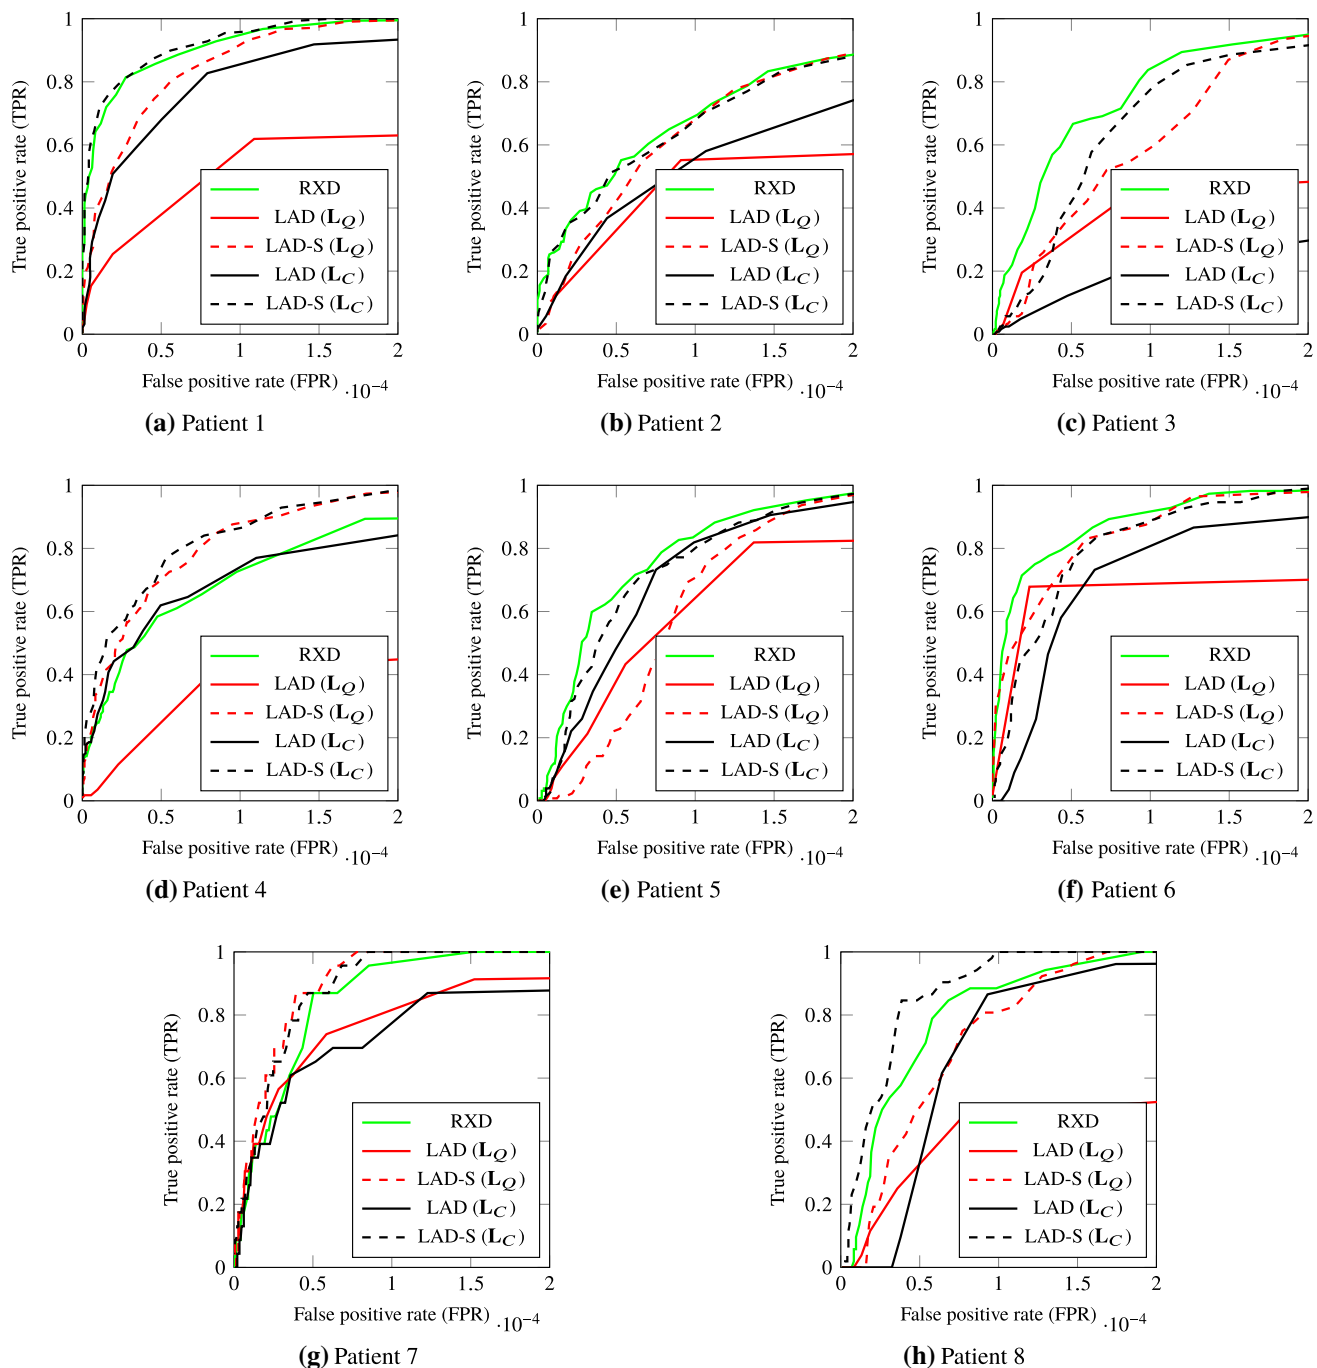


Fig. 8 ROC curves for all patients in the *Tumor* testing scenario

to notice regarding this setup is that we are dealing with voxels and 3D volumes. For this reason, in LAD-S we will use 6-connectivity, which is the extension of 2D 4-connectivity to 3D space.

To compare performance yielded by the different approaches, we use SOI as presented in (18). Once again, in this study we selected the value of t yielding the highest SOI.

Figure 8 shows the ROC curves for all the eight patients in the *Tumor* dataset, while Table 3 shows the average SOI results of our tests over the patient dataset. The inclusion of spatial information in the graph improves the SOI metric. In this scenario, we do not present results after dimensionality reduction because the spectral dimensions were already very few. Also, in this scenario the use of LAD is able to

Table 3 Experimental results in Tumor setup (SOI)

	Average
RXD	0.570
LAD (\mathbf{L}_Q)	0.362
LAD-S (\mathbf{L}_Q)	0.592
LAD (\mathbf{L}_C)	0.427
LAD-S (\mathbf{L}_C)	0.560
Bold values indicates the best result	

obtain performance similar when not better than RXD in all its variances.

6 Discussion

In the previous section, we conducted experiments in hyperspectral and medical domain. RXD's limitations detailed in Sect. 2 can be noticed in many of the presented experiments. In particular, the high number of false negative can be easily noticed in Fig. 3, while the poor performance of RXD, RSAD, and WSCF for the *Impl-4* scenario can be imputed to the fact that in that case the anomaly has a very similar covariance matrix to the background as shown in Fig. 4; this makes very difficult for covariance-based methods to find an acceptable solution.

The results obtained by RSAD have been particularly surprising. The algorithm has been able to achieve results inline or even better than the other two covariance-based approaches in a couple of scenarios, while obtaining very poor performance in the others due to very high FPR. We believe this behavior is caused by the assumption made by RSAD while marking pixels as anomalous that the Mahalanobis distance follows a χ distribution. In the scenarios used in this study, we observed that that was rarely the case. When this assumption does not hold, the decision criterion used by RSAD is probably not sufficient.

The proposed technique was able to outperform state-of-the-art techniques in all scenarios, proving how the flexibility of a graph model can actually enable better and more robust background estimation as well as successful inclusion of spatial information.

Spatially aware variants of the proposed techniques were able to achieve better performance in the *Tumor* scenarios, while failing at improving the performance of the spectral-only variants in the hyperspectral ones. The benefit of including spatial information is more noticeable in the medical scenario because in that case the spectral dimension is reduced to only three bands, representing three different acquisitions in time, as opposed to the 204 spectral bands of the hyperspectral images. Also, we used a uniform correla-

tion as model for the spatial weights; a more refined model might be more suited to better capture the spatial dynamics of remote sensing, while the one used might just be more fitting for medical imaging.

When comparing results obtained by LAD using \mathbf{L}_Q or \mathbf{L}_C , it can be noticed how performance is often very similar on hyperspectral images, while in *Tumor* \mathbf{L}_C is able to obtain consistently better results. This behavior is clearly due to the fact that \mathbf{L}_Q depends on pairwise correlation estimates that are particularly critical in the *Tumor* case, where the 3D volumes are characterized by poor spatiotemporal resolution. In this case, the use of graph prior based on \mathbf{L}_C turns out to be more robust. An analysis of the ROCs validated this observation even further: For the hyperspectral case, the ROC curves for LAD using \mathbf{L}_Q or \mathbf{L}_C behave very similarly in both cases, indicating that the two weight functions are able to capture the same aspects of the data, while in the *Tumor* case, the two ROC curves have a more varied behavior.

All these tests confirm that the use of our approach is preferable to RXD, RSAD, and WSCF and that Laplacian estimated using the Cauchy distance is able to perform as well as the one estimated using partial correlation. Once again, this is remarkable as the former does not require any matrix inversion, while the latter does.

7 Conclusions

We present Laplacian anomaly detector, a graph-based algorithm aiming at detecting targets by virtue of a Laplacian model of the image background. Two different approaches to the graph construction are proposed. When comparing to RXD, RSAD, and WSCF, one of the main advantages of our technique is its ability to model the image content without the need for matrix inversions. Both visual inspection and objective results show how the proposed approach is able to outperform the other benchmark methods. Future direction might be devoted to evaluate LAD ability to detect anomalies on generic non-image graphs.

Acknowledgements Open access funding provided by Aalto University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Acito, N., Diani, M., Corsini, G.: On the CFAR property of the RX algorithm in the presence of signal-dependent noise in hyperspectral images. *IEEE Trans. Geosci. Remote Sens.* **51**(6), 3475–3491 (2013)
- Akoglu, L., Tong, H., Koutra, D.: Graph based anomaly detection and description: a survey. *Data Min. Knowl. Disc.* **29**(3), 626–688 (2014). <https://doi.org/10.1007/s10618-014-0365-y>
- Anandkumar, A., Tong, L., Swami, A.: Detection of Gauss–Markov random fields with nearest-neighbor dependency. *IEEE Trans. Inf. Theory* **55**(2), 816–827 (2009). <https://doi.org/10.1109/TIT.2008.2009855>
- Baghbi, M.Z., Jamshidi, K., Naghsh-Nilchi, A.R., Homayouni, S.: Improvement of anomaly detection algorithms in hyperspectral images using discrete wavelet transform. *Signal Image Process. Int. J.* **2**(4), 13–25 (2011)
- Banerjee, A., Burlina, P., Diehl, C.: A support vector method for anomaly detection in hyperspectral imagery. *IEEE Trans. Geosci. Remote Sens.* **44**(8), 2282–2291 (2006). <https://doi.org/10.1109/TGRS.2006.873019>
- Baterina, A.V., Oppus, C.: Image edge detection using ant colony optimization. *WSEAS Trans. Sig. Proc.* **6**(2), 58–67 (2010)
- Bertozzi, A., Flenner, A.: Diffuse interface models on graphs for classification of high dimensional data. *Multisc. Model. Simul.* **10**(3), 1090–1118 (2012). <https://doi.org/10.1137/11083109X>
- Black, M.J., Sapiro, G., Marimont, D.H., Heeger, D.: Robust anisotropic diffusion. *IEEE Trans. Image Process.* **7**(3), 421–432 (1998). <https://doi.org/10.1109/83.661192>
- Boykov, Y., Funka-Lea, G.: Graph cuts and efficient N–D image segmentation. *Int. J. Comput. Vision* **70**(2), 109–131 (2006). <https://doi.org/10.1007/s11263-006-7934-5>
- Bracco, C., Verdoja, F., Grangetto, M., Di Dia, A., Racca, M., Varetto, T., Stasi, M.: Automatic GTV contouring applying anomaly detection algorithm on dynamic FDG PET images. *Physica Med.* **32**(1), 99 (2016). <https://doi.org/10.1016/j.ejmp.2016.01.343>
- Chalapathy, R., Chawla, S.: Deep learning for anomaly detection: a survey. *arXiv:1901.03407* [cs, stat] (2019)
- Chang, C.I., Chiang, S.S.: Anomaly detection and classification for hyperspectral imagery. *IEEE Trans. Geosci. Remote Sens.* **40**(6), 1314–1325 (2002). <https://doi.org/10.1109/TGRS.2002.800280>
- Chang, C.I., Du, Q.: Noise subspace projection approaches to determination of intrinsic dimensionality of hyperspectral imagery. In: *Proc. SPIE*, vol. 3871, pp. 34–44. Florence, Italy (1999). <https://doi.org/10.1117/12.373271>
- Chang, C.I., Heinz, D.C.: Constrained subpixel target detection for remotely sensed imagery. *IEEE Trans. Geosci. Remote Sens.* **38**(3), 1144–1159 (2000). <https://doi.org/10.1109/36.843007>
- Cheng, G., Zhou, P., Han, J.: Learning rotation-invariant convolutional neural networks for object detection in VHR optical remote sensing images. *IEEE Trans. Geosci. Remote Sens.* **54**(12), 7405–7415 (2016). <https://doi.org/10.1109/TGRS.2016.2601622>
- Cheng, K.W., Chen, Y.T., Fang, W.H.: Gaussian process regression-based video anomaly detection and localization with hierarchical feature representation. *IEEE Trans. Image Process.* **24**(12), 5288–5301 (2015). <https://doi.org/10.1109/TIP.2015.2479561>
- Cheung, G., Magli, E., Tanaka, Y., Ng, M.K.: Graph spectral image processing. *Proc. IEEE* **106**(5), 907–930 (2018). <https://doi.org/10.1109/JPROC.2018.2799702>
- Spectral Graph Theory. No. 92 in Regional conference series in mathematics. American Mathematical Society, Providence (1997)
- Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945). <https://doi.org/10.2307/1932409>
- Du, B., Zhang, L.: Random-selection-based anomaly detector for hyperspectral imagery. *IEEE Trans. Geosci. Remote Sens.* **49**(5), 1578–1589 (2011). <https://doi.org/10.1109/TGRS.2010.2081677>
- Du, B., Zhang, L.: A discriminative metric learning based anomaly detection method. *IEEE Trans. Geosci. Remote Sens.* **52**(11), 6844–6857 (2014). <https://doi.org/10.1109/TGRS.2014.2303895>
- Du, B., Zhang, L.: Target detection based on a dynamic subspace. *Pattern Recogn.* **47**(1), 344–358 (2014). <https://doi.org/10.1016/j.patcog.2013.07.005>
- Du, B., Zhang, Y., Zhang, L., Tao, D.: Beyond the sparsity-based target detector: a hybrid sparsity and statistics-based detector for hyperspectral images. *IEEE Trans. Image Process.* **25**(11), 5345–5357 (2016). <https://doi.org/10.1109/TIP.2016.2601268>
- Du, Q., Ren, H.: Real-time constrained linear discriminant analysis to target detection and classification in hyperspectral imagery. *Pattern Recogn.* **36**(1), 1–12 (2003). [https://doi.org/10.1016/S0031-3203\(02\)00065-1](https://doi.org/10.1016/S0031-3203(02)00065-1)
- Egilmez, H.E., Ortega, A.: Spectral anomaly detection using graph-based filtering for wireless sensor networks. In: 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1085–1089 (2014)
- Fowler, J.E., Du, Q.: Anomaly detection and reconstruction from random projections. *IEEE Trans. Image Process.* **21**(1), 184–195 (2012). <https://doi.org/10.1109/TIP.2011.2159730>
- Fracastoro, G., Fosson, S.M., Magli, E.: Steerable discrete cosine transform. *IEEE Trans. Image Process.* **26**(1), 303–314 (2017). <https://doi.org/10.1109/TIP.2016.2623489>
- Fracastoro, G., Magli, E.: Predictive graph construction for image compression. In: IEEE International Conference on Image Processing (ICIP), pp. 2204–2208 (2015). <https://doi.org/10.1109/ICIP.2015.7351192>
- Galasso, F., Keuper, M., Brox, T., Schiele, B.: Spectral graph reduction for efficient image and streaming video segmentation. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4321–4328 (2014)
- Garber, K.: Energy boost: the Warburg effect returns in a new theory of cancer. *JNCI J. Natl. Cancer Inst.* **96**(24), 1805–1806 (2004). <https://doi.org/10.1093/jnci/96.24.1805>
- Gaucel, J.M., Guillaume, M., Bourennane, S.: Whitening spatial correlation filtering for hyperspectral anomaly detection. In: Proceedings of the (ICASSP '05) IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005, vol. 5, pp. v/333–v/336 (2005). <https://doi.org/10.1109/ICASSP.2005.1416308>
- Grady, L.J., Polimeni, J.R.: Discrete Calculus: Applied Analysis on Graphs for Computational Science. Springer, London (2010)
- Gu, Y., Liu, Y., Zhang, Y.: A selective KPCA algorithm based on high-order statistics for anomaly detection in hyperspectral imagery. *IEEE Geosci. Remote Sens. Lett.* **5**(1), 43–47 (2008). <https://doi.org/10.1109/LGRS.2007.907304>
- Gurram, P., Kwon, H.: Support-vector-based hyperspectral anomaly detection using optimized kernel parameters. *IEEE Geosci. Remote Sens. Lett.* **8**(6), 1060–1064 (2011). <https://doi.org/10.1109/LGRS.2011.2155030>
- Han, J., Ji, X., Hu, X., Zhu, D., Li, K., Jiang, X., Cui, G., Guo, L., Liu, T.: Representing and retrieving video shots in human-centric brain imaging space. *IEEE Trans. Image Process.* **22**(7), 2723–2736 (2013). <https://doi.org/10.1109/TIP.2013.2256919>
- Hanczar, B., Hua, J., Sima, C., Weinstein, J., Bittner, M., Dougherty, E.R.: Small-sample precision of ROC-related estimates. *Bioinformatics* **26**(6), 822–830 (2010). <https://doi.org/10.1093/bioinformatics/btq037>
- Hand, D.J.: Measuring classifier performance: a coherent alternative to the area under the ROC curve. *Mach. Learn.* **77**(1), 103–123 (2009). <https://doi.org/10.1007/s10994-009-5119-5>
- Harsanyi, J.C., Farrand, W.H., Chang, C.I.: Determining the number and identity of spectral endmembers: an integrated approach

- using Neyman–Pearson Eigen–thresholding and iterative constrained RMS error minimization. In: *Proceedings of the Thematic Conference on Geologic Remote Sensing*, vol. 1, pp. 395–395. Environmental Research Institute of Michigan (1993)
39. Hu, W., Cheung, G., Ortega, A., Au, O.C.: Multiresolution graph fourier transform for compression of piecewise smooth images. *IEEE Trans. Image Process.* **24**(1), 419–433 (2015). <https://doi.org/10.1109/TIP.2014.2378055>
 40. Khazai, S., Homayouni, S., Safari, A., Mojaradi, B.: Anomaly detection in hyperspectral images based on an adaptive support vector method. *IEEE Geosci. Remote Sens. Lett.* **8**(4), 646–650 (2011). <https://doi.org/10.1109/LGRS.2010.2098842>
 41. Kwon, H., Nasrabadi, N.M.: Kernel RX-algorithm: a nonlinear anomaly detector for hyperspectral imagery. *IEEE Trans. Geosci. Remote Sens.* **43**(2), 388–397 (2005). <https://doi.org/10.1109/TGRS.2004.841487>
 42. Kwon, H., Nasrabadi, N.M.: Kernel matched subspace detectors for hyperspectral target detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(2), 178–194 (2006). <https://doi.org/10.1109/TPAMI.2006.39>
 43. Li, W., Du, Q.: Collaborative representation for hyperspectral anomaly detection. *IEEE Trans. Geosci. Remote Sens.* **53**(3), 1463–1474 (2015). <https://doi.org/10.1109/TGRS.2014.2343955>
 44. Li, W., Wu, G., Du, Q.: Transferred deep learning for anomaly detection in hyperspectral imagery. *IEEE Geosci. Remote Sens. Lett.* **14**(5), 597–601 (2017). <https://doi.org/10.1109/LGRS.2017.2657818>
 45. Lobo, J.M., Jiménez-Valverde, A., Real, R.: AUC: a misleading measure of the performance of predictive distribution models. *Glob. Ecol. Biogeogr.* **17**(2), 145–151 (2008). <https://doi.org/10.1111/j.1466-8238.2007.00358.x>
 46. Lézoray, O., Grady, L.: *Image Processing and Analysis with Graphs: Theory and Practice*. CRC Press, Boca Raton (2012)
 47. Mahalanobis, P.C.: On the generalized distance in statistics. In: *National Institute of Sciences of India*, vol. 2, pp. 49–55. Calcutta, India (1936)
 48. Manolakis, D.G., Lockwood, R., Cooley, T., Jacobson, J.: Is there a best hyperspectral detection algorithm? In: *Proceedings of the SPIE*, vol. 7334, pp. 733402–733402–16 (2009). <https://doi.org/10.1117/12.816917>
 49. Matteoli, S., Diani, M., Corsini, G.: A tutorial overview of anomaly detection in hyperspectral images. *IEEE Aerosp. Electron. Syst. Mag.* **25**(7), 5–28 (2010). <https://doi.org/10.1109/MAES.2010.5546306>
 50. Matteoli, S., Diani, M., Corsini, G.: Hyperspectral anomaly detection with Kurtosis–Driven local covariance matrix corruption mitigation. *IEEE Geosci. Remote Sens. Lett.* **8**(3), 532–536 (2011). <https://doi.org/10.1109/LGRS.2010.2090337>
 51. Matteoli, S., Veracini, T., Diani, M., Corsini, G.: Models and methods for automated background density estimation in hyperspectral anomaly detection. *IEEE Trans. Geosci. Remote Sens.* **51**(5), 2837–2852 (2013)
 52. Noble, C.C., Cook, D.J.: Graph-based Anomaly Detection. In: *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '03*, pp. 631–636. ACM, New York, NY, USA (2003). <https://doi.org/10.1145/956750.956831>
 53. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Trans. Systems Man Cybern.* **9**(1), 62–66 (1979). <https://doi.org/10.1109/TSMC.1979.4310076>
 54. Perez, C.A., Brady, L.W.: *Principles and Practice of Radiation Oncology*, 5th edn. Lippincott Williams & Wilkins, Philadelphia, PA (2008)
 55. Ravazzi, C., Coluccia, G., Magli, E.: Curl-constrained gradient estimation for image recovery from highly incomplete spectral data. *IEEE Trans. Image Process.* **PP**(99):1–1 (2017)
 56. Reed, I.S., Yu, X.: Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution. *IEEE Trans. Acoust. Speech Signal Process.* **38**(10), 1760–1770 (1990). <https://doi.org/10.1109/29.60107>
 57. Sandryhaila, A., Moura, J.M.F.: Discrete signal processing on graphs: graph fourier transform. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6167–6170 (2013). <https://doi.org/10.1109/ICASSP.2013.6638850>
 58. Sandryhaila, A., Moura, J.M.F.: Discrete signal processing on graphs: frequency analysis. *IEEE Trans. Signal Process.* **62**(12), 3042–3054 (2014). <https://doi.org/10.1109/TSP.2014.2321121>
 59. Santner, J., Pock, T., Bischof, H.: Interactive multi-label segmentation. In: *Computer Vision—ACCV 2010, Lecture Notes in Computer Science*, vol. 6492, pp. 397–410. Springer, Berlin (2011)
 60. Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P.: The emerging field of signal processing on graphs: extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process. Mag.* **30**(3), 83–98 (2013). <https://doi.org/10.1109/MSP.2012.2235192>
 61. Stefanou, M.S., Kerekes, J.P.: A method for assessing spectral image utility. *IEEE Trans. Geosci. Remote Sens.* **47**(6), 1698–1706 (2009). <https://doi.org/10.1109/TGRS.2008.2006364>
 62. Stein, D.W.J., Beaven, S.G., Hoff, L.E., Winter, E.M., Schaum, A.P., Stocker, A.D.: Anomaly detection from hyperspectral imagery. *IEEE Signal Process. Mag.* **19**(1), 58–69 (2002). <https://doi.org/10.1109/79.974730>
 63. Verdoja, F., Bonafè, B., Cavagnino, D., Grangetto, M., Bracco, C., Varetto, T., Racca, M., Stasi, M.: Global and local anomaly detectors for tumor segmentation in dynamic PET acquisitions. In: *2016 IEEE International Conference on Image Processing (ICIP)*, pp. 4131–4135. IEEE, Phoenix, AZ (2016). <https://doi.org/10.1109/ICIP.2016.7533137>
 64. Verdoja, F., Grangetto, M.: Directional graph weight prediction for image compression. In: *IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP 2017)*, pp. 1517–1521. IEEE, New Orleans, LA (2017)
 65. Verdoja, F., Grangetto, M., Bracco, C., Varetto, T., Racca, M., Stasi, M.: Automatic method for tumor segmentation from 3-points dynamic PET acquisitions. In: *IEEE International Conference on Image Processing 2014 (ICIP 2014)*, pp. 937–941. IEEE, Paris, France (2014). <https://doi.org/10.1109/ICIP.2014.7025188>
 66. Wong, K.P., Feng, D., Meikle, S.R., Fulham, M.J.: Segmentation of dynamic PET images using cluster analysis. *IEEE Trans. Nucl. Sci.* **49**(1), 200–207 (2002). <https://doi.org/10.1109/TNS.2002.998752>
 67. Yuan, Y., Wang, Q., Zhu, G.: Fast hyperspectral anomaly detection via high-order 2-D crossing filter. *IEEE Trans. Geosci. Remote Sens.* **53**(2), 620–630 (2015). <https://doi.org/10.1109/TGRS.2014.2326654>
 68. Yuan, Z., Sun, H., Ji, K., Li, Z., Zou, H.: Local sparsity divergence for hyperspectral anomaly detection. *IEEE Geosci. Remote Sens. Lett.* **11**(10), 1697–1701 (2014). <https://doi.org/10.1109/LGRS.2014.2306209>
 69. Zaidi, H., Abdoli, M., Fuentes, C.L., El Naqa, I.M.: Comparative methods for PET image segmentation in pharyngolaryngeal squamous cell carcinoma. *Eur. J. Nucl. Med. Mol. Imaging* **39**(5), 881–891 (2012). <https://doi.org/10.1007/s00259-011-2053-0>
 70. Zhang, C., Florêncio, D.: Analyzing the optimality of predictive transform coding using graph-based models. *IEEE Signal Processing Letters* (2013)
 71. Zhang, D., Meng, D., Han, J.: Co-saliency detection via a self-paced multiple-instance learning framework. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(5), 865–878 (2017). <https://doi.org/10.1109/TPAMI.2016.2567393>
 72. Zhang, Y., Du, B., Zhang, L.: A sparse representation-based binary hypothesis model for target detection in hyperspectral images.

- IEEE Trans. Geosci. Remote Sens. **53**(3), 1346–1354 (2015). <https://doi.org/10.1109/TGRS.2014.2337883>
73. Zou, K.H., Warfield, S.K., Bharatha, A., Tempany, C.M.C., Kaus, M.R., Haker, S.J., Wells, W.M., Jolesz, F.A., Kikinis, R.: Statistical validation of image segmentation quality based on a spatial overlap index. *Academic Radiology* **11**(2), 178–189 (2004). [https://doi.org/10.1016/S1076-6332\(03\)00671-8](https://doi.org/10.1016/S1076-6332(03)00671-8)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Francesco Verdoja received his PhD in Computer Science from University of Turin, Italy, in 2017. During his PhD studies, his research focused on graph-based image processing. He has also conducted research on image and 3D point cloud segmentation, and his work on tumor segmentation awarded him the “Best student award” at the International Computer Vision Summer School 2014 (ICVSS14). Since September 2017, he joined the Intelligent Robotics group at Aalto University in Helsinki, Finland, as a postdoctoral researcher. His research now focuses on robotic perception and mapping.

Marco Grangetto received the Electrical Engineering and PhD degrees from the Politecnico di Torino, Torino, Italy, in 1999 and 2003, respectively. He is currently a Full Professor at the Computer Science Department, Università di Torino. His research interests are in the fields of digital communications, multimedia signal processing and computer vision. In particular, his expertise includes data compression, error resilient video coding, joint source channel coding, distributed source coding, image and video processing. Prof. Grangetto is an Associate Editor of the IEEE Transactions on Multimedia. He was the recipient of a Fulbright grant in 2001 for a research period with the Department of Electrical and Computer Engineering, University of California at San Diego, San Diego, CA, USA.