



HAL
open science

Coupled dictionary learning for unsupervised change detection between multimodal remote sensing images

Vinicius Ferraris, Nicolas Dobigeon, Yanna Cruz Cavalcanti, Thomas Oberlin,
Marie Chabert

► **To cite this version:**

Vinicius Ferraris, Nicolas Dobigeon, Yanna Cruz Cavalcanti, Thomas Oberlin, Marie Chabert. Coupled dictionary learning for unsupervised change detection between multimodal remote sensing images. Computer Vision and Image Understanding, 2019, 189, pp.1-15. 10.1016/j.cviu.2019.102817. hal-02397250

HAL Id: hal-02397250

<https://hal.science/hal-02397250>

Submitted on 6 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:

<http://oatao.univ-toulouse.fr/25034>

Official URL

DOI : <https://doi.org/10.1016/j.cviu.2019.102817>

To cite this version: Ferraris, Vinicius and Dobigeon, Nicolas and Cruz Cavalcanti, Yanna and Oberlin, Thomas and Chabert, Marie *Coupled dictionary learning for unsupervised change detection between multimodal remote sensing images.* (2019) *Computer Vision and Image Understanding*, 189. 1-15. ISSN 1077-3142

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Coupled dictionary learning for unsupervised change detection between multimodal remote sensing images^{☆,☆☆}

Vinicius Ferraris^{*}, Nicolas Dobigeon, Yanna Cavalcanti, Thomas Oberlin, Marie Chabert

University of Toulouse, IRIT/INP-ENSEEIH, 2 Rue Camichel, 31071 Toulouse, France

A B S T R A C T

Keywords:

Change detection
Coupled dictionary learning
Multimodality
Optical images
SAR image

Archetypal scenarios for change detection generally consider two images acquired through sensors of the same modality. However, in some specific cases such as emergency situations, the only images available may be those acquired through sensors of different modalities. This paper addresses the problem of unsupervisedly detecting changes between two observed images acquired by sensors of different modalities with possibly different resolutions. These sensor dissimilarities introduce additional issues in the context of operational change detection that are not addressed by most of the classical methods. This paper introduces a novel framework to effectively exploit the available information by modeling the two observed images as a sparse linear combination of atoms belonging to a pair of coupled overcomplete dictionaries learnt from each observed image. As they cover the same geographical location, codes are expected to be globally similar, except for possible changes in sparse spatial locations. Thus, the change detection task is envisioned through a dual code estimation which enforces spatial sparsity in the difference between the estimated codes associated with each image. This problem is formulated as an inverse problem which is iteratively solved using an efficient proximal alternating minimization algorithm accounting for nonsmooth and nonconvex functions. The proposed method is applied to real images with simulated yet realistic and real changes. A comparison with state-of-the-art change detection methods evidences the accuracy of the proposed strategy.

1. Introduction

Ecosystems exhibit permanent variations at different temporal and spatial scales caused by natural, anthropogenic, or even both factors (Coppin et al., 2004). Monitoring spatial variations over a period of time is an important source of knowledge that helps understanding the possible transformations occurring on Earth's surface. Therefore, due to the importance of quantifying these transformations, change detection (CD) has been an ubiquitous issue addressed in the remote sensing and geoscience literature (Lu et al., 2004).

Remote sensing CD methods can be first classified with respect to (w.r.t.) their supervision (Bovolo and Bruzzone, 2007), depending on the availability of prior knowledge about the expected changes. More precisely, supervised CD methods require ground reference information about at least one of the observations. Conversely, unsupervised CD can be contextualized as automatic detection of changes without the need for any further external knowledge. Each class of CD methods present particular competitive advantages w.r.t. the others. For instance, supervised CD methods generally achieve better accuracy for predefined

modalities whereas unsupervised methods are characterized by their flexibility and genericity. Nevertheless, implementing supervised methods require the acquisition of relevant ground information, which is a very challenging and expensive task, in terms of human and time resources (Bovolo and Bruzzone, 2007). Relaxing this constraints makes unsupervised methods more suitable for operational CD.

CD methods can also be categorized w.r.t. the imagery modalities the method is able to handle. As remote sensing encompasses many different types of imagery modalities (e.g., single- and multi-band optical images, radar, LiDAR), dedicated CD methods have been specifically developed for each one by exploiting its acquisition process and the intrinsic characteristics of the resulting data. Thus, due to differences in the physical meaning and statistical properties of images from different sensor modalities, a general CD method able to handle all modalities is particularly difficult to design and to implement. For this reason, most of the CD methods focus on a pair of images from one single target modality. In this case, the images are generally compared pixel-wisely using the underlying assumption of same spatial resolutions (Singh,

[☆] Part of this work has been supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), Brazil, the EU FP7 through the ERANETMED JC-WATER program [MapInvPlnt Project ANR-15-NMED-0002-02] and the ANR, France-3IA Artificial and Natural Intelligence Toulouse Institute (ANITI).

^{☆☆} No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to <https://doi.org/10.1016/j.cviu.2019.102817>.

^{*} Corresponding author.

E-mail address: vinicius.ferraris@enseiht.fr (V. Ferraris).

1989; Bovolo and Bruzzone, 2015). Nevertheless, in some practical scenarios such as, e.g., emergency missions due to natural disasters, when the availability of data and the responsiveness are strong constraints, CD methods may have to handle observations of different modalities and/or resolutions. This highlights the need for robust and flexible CD techniques able to deal with this kind of observations.

The literature about multimodal CD is very limited, yet a few relevant references include the works by Kawamura (1971), Bruzzone et al. (1999), Inglada (2002), Lu et al. (2004), Alberga et al. (2007), Mercier et al. (2008) and Prendes et al. (2015). However, multimodal CD has always been an important topic since the initial development of CD methods. Earlier work by Kawamura (1971) described the potential of CD between a multimodal collection of datasets (e.g., photographic, infrared and radar), applied to weather prediction and land surveillance. Three features are extracted from the pair and the CD algorithm is trained on a learning set. According to Lu et al. (2004), various methods dedicated to CD between images from different sources of data are grouped as geographical information system-based methods. For instance, Solberg et al. (1996) proposed a supervised classification of multisource satellite images using Markov random fields. The work of Bruzzone et al. (1999) uses compound classification to detect changes in multisource data. The method uses artificial neural networks to estimate the posterior probability of classes. Moreover, Inglada (2002) studies the relevance of several similarity measures between multisensor data. These measures are implemented in a CD context (Alberga et al., 2007). A preprocessing technique based on conditional copula that contributes to better statistically modeling multisensor images was proposed by Mercier et al. (2008). Besides, Brunner et al. (2010) presented a strategy to assess building damages using a pair of very high resolution (VHR) optical and radar images by geometrically modeling buildings in both modalities. Chabert et al. (2010) updated information databases by means of logistic regression. More recently, the work of Prendes et al. (2015) presented a supervised method to infer changes after learning a manifold defined by pairs of patches extracted from the two images. Although some of these methods present relatively high detection performance, they are often restrained to a specific modality or to a specific target application. For instance, Solano-Correa et al. (2018) proposed an approach to detect changes between multispectral images with different spatial and spectral resolutions by homogenization of radiometric and geometric image properties. However, this approach relies in particular on a taxonomy of possible radiometric changes observed in very high resolution images. Moreover, some methods are only suitable for building damage assessment taking benefit of their high-level modeling, but show a poor adaptability to other scenarios (Brunner et al., 2010; Chabert et al., 2010). The other ones estimate some metrics from unchanged trained samples, which prevents their application within a fully unsupervised context (Bruzzone et al., 1999; Prendes et al., 2015; Mercier et al., 2008).

Recently, an unsupervised multi-source CD method based on coupled dictionary learning was addressed by Gong et al. (2016). In the proposed methodology, the CD is based on the reconstruction error of patches approximated thanks to estimated coupled dictionary and independent sparse codes. Atoms of the dictionary are learnt from pairs of patches jointly extracted from the observed images. Following the same principle, in Lu et al. (2017), a semi-supervised method was used to handle multispectral images based on joint dictionary learning. Both methods rely on the rationale that the coupled dictionary estimated from the observed images tends to produce stronger reconstruction errors in change regions rather than in unchanged ones. Because of the multi-modality, the problem has not been formulated in the image space, but rather in a latent space formed by the coupled dictionary atoms. However, both methods exhibit some crucial issues that may impair their relative performance. First, the underlying optimization problem is highly nonconvex and no convergence guarantees are ensured, even by using some traditional dictionary learning methods (Aharon et al., 2006). Then, the considered CD problem has been

split into two distinct steps: dictionary learning and code estimation. The errors in code estimation may produce false alarms in the final CD even with reliable dictionary estimates. Also, the statistical model of the noise – inherent to each sensor modality – has not been taken into consideration explicitly, which may dramatically impact the CD performance (Campbell and Wynne, 2011). Finally, these methods do not consider overlapping patches, which potentially would increase their robustness and thus do not explicitly handle the problem of possible differences in spatial resolutions (Ferraris et al., 2017b,a). The adequacy between the size of patches and the image scale is not discussed, although it may have a negative impact on the dictionary coupling and thus on the detection performance.

Overcoming these limitations, this paper proposes a similar methodology to learn coupled dictionaries able to conveniently model multimodal remote sensing images. Specifically, contrary to the aforementioned methods, the problem is fully formulated without splitting the learning and coding steps. Also, an appropriate statistical model is derived to describe the image from each specific remote sensing modality. Besides, the proposed method explicitly allows patch overlapping within the overall estimation process. To couple images with different resolutions, additional scaling matrices inspired by the work by Seichepine et al. (2014) are jointly estimated within the whole process. Finally, as the problem is highly nonconvex, it is iteratively solved based on the proximal alternating linearized minimization (PALM) algorithm (Bolte et al., 2014), which ensures convergence towards a critical point for some nonconvex nonsmooth problems. Note that the proposed patch-based method departs from segmentation-based methods which generally extract change information at an object-level, whose resolution is implicitly defined by the chosen segmentation procedure (Feng et al., 2018). Instead, the proposed method linearly decomposes overlapping square patches onto an appropriate common latent space, which allows CD to be operated at a pixel level.

This manuscript is organized as follows. Generic and well-admitted image models are introduced in Section 2. Capitalizing on these image models, Section 3 formulated the CD problem as a coupled dictionary learning. Section 4 proposes an algorithmic solution to minimize the resulting CD-based objective function. Section 5 reports experimental results obtained on synthetic images, considering three distinct simulation scenarios. Experiments conducted on real images are presented in Section 5.3. Finally, Section 6 concludes the manuscript.

2. Image models

2.1. Forward model

Let us consider that the image formation process inherent to all digital remote sensing imagery modalities is modeled as a sequence of transformations, denoted $T[\cdot]$. This sequence applies to the original scene to produce the sensor output image. This output image is referred to as the observed image and is denoted by $\mathbf{Y} \in \mathbb{R}^{L \times N}$ consisting of N voxels $\mathbf{y}_i \in \mathbb{R}^L$ stacked lexicographically that is from left to right, row by row. The voxel dimension L may represent different quantities depending on the modality of the data. For instance, it stands for the number of spectral bands in the case of multiband optical images (Ferraris et al., 2017b) or for the number of polarization modes in the case of polarimetric synthetic aperture radar (POLSAR) images. The observed image provides a limited representation of the original scene with properties imposed by the image signal processor characterizing the sensor. The original scene cannot be exactly represented because of its continuous nature, but it can be conveniently approximated by a latent (i.e., unobserved) image $\mathbf{X} \in \mathbb{R}^{L \times N}$ related to the observed image as follows

$$\mathbf{Y} = T[\mathbf{X}]. \quad (1)$$

The sequence of transformations $T[\cdot]$ operated by the sensor over the latent image is often referred to as the *degradation process*. It may

represent resolution degradations accounting for the spatial and/or spectral characteristics of the sensor (Ferraris et al., 2017a,b). In this paper, it specifically models the intrinsic noise corruption associated to the sensor modality (Sun and Févotte, 2014). The latent image \mathbf{X} can be understood, in this context, as a noise-free version of the observed image \mathbf{Y} with the same resolution.

More precisely, the transformation $T[\cdot]$ underlies the likelihood function $p(\mathbf{Y}|\mathbf{X})$ which statistically models the observed image \mathbf{Y} conditionally to the latent image \mathbf{X} by taking into account the noise statistics. The noise statistical model mainly depends on the modality and rely on some classical distributions, e.g., the Gaussian distribution for optical images or the Gamma distribution for multi-look SAR images. Moreover, as already pointed out by Févotte et al. (2009) in a different application context, for a wide family of distributions, this likelihood function relies on a divergence measure $D(\cdot|\cdot)$ between the observed and latent images, which finally defines an explicit data-fitting term through a negative-log transformation

$$-\log p(\mathbf{Y}|\mathbf{X}) = \phi^{-1}D(\mathbf{Y}|\mathbf{X}) + \theta \quad (2)$$

where ϕ and θ are parameters characterizing the distributions. In Appendix A, the divergence measures $D(\cdot|\cdot)$ are derived for two of the most common remote sensing image modalities, namely optical multiband and SAR images, considered in this work.

2.2. Latent image sparse model

Sparse representations have been an ubiquitous and well-admitted tool to model images in various applications and task-driven contexts (Mairal, 2014). Indeed, natural images are known to be compressible in a transformed domain, i.e., they can be efficiently represented by a few expansion coefficients acting on basis functions (Mallat, 2009). This finding has motivated numerous works on image understanding, compression and denoising (Olshausen and Field, 1997; Chen et al., 2001). In earlier works, this transformed domain, equivalently defined by the associated basis functions, was generally fixed in advance and chosen in agreement with the expected spatial content of the images (Mallat, 2009). Thus, the basis functions belonged to pre-determined families with specific representation abilities, such as cosines, wavelets, contourlets, shearlets, among others. More recently, the seminal contribution by Aharon et al. proposed a new paradigm by learning an overcomplete dictionary jointly with a sparse code (Aharon et al., 2006). This dictionary learning-based approach exploits the key property of self-similarity characterizing the images to provide an adaptive representation. Indeed, it aims at identifying elementary patches that can be linearly and sparsely combined to approximate the observed image patches. In this paper, following the approach by Aharon et al. (2006), we propose to resort to this dictionary-based representation to model the latent image \mathbf{X} . More precisely, the image is first decomposed into a set of N_p 3D-patches with $1 \leq N_p \leq N$. Let $\mathcal{R}_i : \mathbb{R}^{L \times N} \rightarrow \mathbb{R}^{K^2 L}$ denote a binary operation modeling the extraction, from the image, of the i th patch ($i \in \{1, \dots, N_p\}$) such that

$$\mathbf{p}_i = \mathcal{R}_i \mathbf{X} \quad (3)$$

where $\mathbf{p}_i \in \mathbb{R}^{K^2 L}$ stands for the i th $K \times K \times L$ -pixel patch in its vectorized form. The integer $K > 1$ defines the spatial size of the patches, i.e., its number of rows and columns before being vectorized. Note that the number of patches N_p is such that $1 \leq N_p \leq N$ and patches may overlap. The choice of the number N_p of patches will be more deeply discussed in Section 3 in the specific context of CD. The conjugate of the patch-extraction operator,¹ denoted \mathcal{R}_i^T , acts on \mathbf{p}_i to produce a zero-padded image composed by the unique patch \mathbf{p}_i located at the i th spatial position.

¹ Note that, despite a slight abuse of notation, the operator \mathcal{R} (resp., \mathcal{R}^T) does not stand for a matrix, but rather for a linear operator acting on the image \mathbf{X} (resp., the patch \mathbf{p}_i) directly.

In accordance with dictionary-based representation principles, these patches are assumed to be approximately and independently modeled as sparse combinations of atoms belonging to an overcomplete dictionary $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_{N_d}] \in \mathbb{R}^{K^2 L \times N_d}$

$$\mathbf{p}_i | \mathbf{D}, \mathbf{a}_i \sim \mathcal{N}(\mathbf{D} \mathbf{a}_i, \sigma^2 \mathbf{I}_{N_d}) \quad (4)$$

where $N_d > 0$ stands for the user-defined number of atoms composing the dictionary, commonly referred to as dictionary size and $\mathbf{a}_i \in \mathbb{R}^{N_d}$ represents the code (coefficients) of the current patch over the dictionary, $\Sigma = \sigma^2 \mathbf{I}_{N_d}$ is the error covariance matrix and \cdot . Let $\mathbf{P} \in \mathbb{R}^{K^2 L \times N_p} = [\mathbf{p}_1, \dots, \mathbf{p}_{N_p}]$ denote the matrix that stacks the set of all, possibly overlapping, patches extracted from the latent image \mathbf{X} at N_p spatial positions arranged on a generally regular spatial grid and enumerated in a lexicographical order (i.e., from left to right and top to bottom of the image). The matrix $\mathbf{A} \in \mathbb{R}^{N_d \times N_p} = [\mathbf{a}_1, \dots, \mathbf{a}_{N_p}]$ is the code matrix in which each column represents the code for each corresponding column of \mathbf{P} . The overcompleteness property of the dictionary, occurring when the number of atoms is greater than the effective dimensionality of the input space, $N_d \gg K^2 L$, allows for the sparsity of the representation (Olshausen and Field, 1997). The overcompleteness implies redundancy and non-orthogonality between atoms. This property is not necessary for the decomposition, but has been proved to be very useful in some applications like denoising and compression (Aharon et al., 2006). Given the image patch matrix \mathbf{P} , dictionary learning aims at recovering the set of atoms \mathbf{D} and the associated code matrix \mathbf{A} and it is generally tackled through a 2-step procedure. First, inferring the code matrix \mathbf{A} associated with the patch matrix \mathbf{P} and the dictionary \mathbf{D} can be formulated as a set of N_p sparsity-penalized optimization problems. Sparsity of the code vectors $\mathbf{a}_i = [a_{i1}, \dots, a_{iN_d}]^T$ ($i = 1, \dots, N_p$) can be promoted by minimizing its ℓ_0 -norm. However, since this leads to a non-convex problem (Chen et al., 2001), it is generally substituted by the corresponding convex relaxation, i.e., an ℓ_1 -norm. Within a probabilistic framework, taking into account the expected non-negativity of the code, this choice can be formulated by assigning a single-side exponential (i.e., Laplacian) prior distribution to the code components, assumed to be a priori independent

$$\mathbf{a}_i \sim \prod_{j=1}^{N_d} \mathcal{L}(a_{ij}; \lambda) \quad (5)$$

where λ is the hyperparameter adjusting the sparsity level over the code. Conversely, learning the dictionary \mathbf{D} given the code \mathbf{A} can also be formulated as an optimization problem. As the number of solutions for the dictionary learning problem can be extremely large, one common assumption is to constrain the energy of each atom, thereby preventing \mathbf{D} to become arbitrarily large (Mairal et al., 2009). Moreover, in the particular context considered in this work, to promote the positivity of the reconstructed patches, the atoms are also constrained to positive values. Thus, each atom will be constrained to the set

$$\mathcal{S} \triangleq \left\{ \mathbf{D} \in \mathbb{R}_+^{K^2 L \times N_d} \mid \forall j \in \{1, \dots, N_d\}, \|\mathbf{d}_j\|_2 = 1 \right\}. \quad (6)$$

2.3. Optimization problem

Adopting a Bayesian probabilistic formulation of the image model introduced in Sections 2.1 and 2.2, the posterior probability of the unknown variables \mathbf{X} , \mathbf{D} and \mathbf{A} can be derived using the probability chain rule (Gelman, 2004)

$$p(\mathbf{X}, \mathbf{D}, \mathbf{A} | \mathbf{Y}) \propto p(\mathbf{Y} | \mathbf{X}) p(\mathbf{X} | \mathbf{D}, \mathbf{A}) p(\mathbf{D}) p(\mathbf{A}) \quad (7)$$

where $p(\mathbf{Y} | \mathbf{X})$ is the likelihood function relating the observation data to the latent image through the direct model (1), $p(\mathbf{X} | \mathbf{D}, \mathbf{A})$ is the dictionary-based prior model of the latent image, $p(\mathbf{D})$ and $p(\mathbf{A})$ are

the (hyper-)prior distributions associated with the dictionary and the sparse code. Under a maximum a posteriori (MAP) paradigm, the joint MAP estimator $\{\hat{\mathbf{X}}_{\text{MAP}}, \hat{\mathbf{D}}_{\text{MAP}}, \hat{\mathbf{A}}_{\text{MAP}}\}$ can be derived by minimizing the negative log-posterior, leading to the following minimization problem

$$\{\hat{\mathbf{X}}_{\text{MAP}}, \hat{\mathbf{D}}_{\text{MAP}}, \hat{\mathbf{A}}_{\text{MAP}}\} \in \operatorname{argmin}_{\mathbf{X}, \mathbf{D}, \mathbf{A}} \mathcal{J}(\mathbf{X}, \mathbf{D}, \mathbf{A}) \quad (8)$$

with

$$\begin{aligned} \mathcal{J}(\mathbf{X}, \mathbf{D}, \mathbf{A}) = & D(\mathbf{Y}|\mathbf{X}) \\ & + \frac{\sigma^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_i \mathbf{X} - \mathbf{D} \mathbf{a}_i\|_F^2 + \\ & + \lambda \|\mathbf{A}\|_1 + \iota_S(\mathbf{D}) \end{aligned} \quad (9)$$

where ι_S represents the indicator function on the set S ,

$$\iota_S(z) = \begin{cases} 0 & \text{if } z \in S \\ +\infty & \text{elsewhere} \end{cases} \quad (10)$$

and $D(\cdot)$ is the data-fitting term associated with the image modality.

This model has been widely advocated in the literature, e.g., for denoising images of various modalities (Elad and Aharon, 2006; Ma et al., 2013). Particularly, in Ma et al. (2013), an additional regularization $\Psi(\mathbf{X})$ of the latent image was introduced as the target modalities may present strong fluctuations due to their inherent image formation process, i.e. Poissonian or multiplicative gamma processes. The final objective function (9) can thus be rewritten as

$$\begin{aligned} \mathcal{J}(\mathbf{X}, \mathbf{D}, \mathbf{A}) = & D(\mathbf{Y}|\mathbf{X}) \\ & + \frac{\sigma^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_i \mathbf{X} - \mathbf{D} \mathbf{a}_i\|_F^2 + \Psi(\mathbf{X}) \\ & + \lambda \|\mathbf{A}\|_1 + \iota_S(\mathbf{D}) \end{aligned} \quad (11)$$

where, for instance, $\Psi(\mathbf{X})$ can stand for a total-variation (TV) regularization (Ma et al., 2013).

The next section expands the proposed image models to handle a pair of observed images in the specific context of CD.

3. From change detection to coupled dictionary learning

3.1. Problem statement

Let us consider two geographically aligned observed images $\mathbf{Y}_1 \in \mathbb{R}^{L_1 \times N_1}$ and $\mathbf{Y}_2 \in \mathbb{R}^{L_2 \times N_2}$ acquired by two sensors S_1 and S_2 at times t_1 and t_2 , respectively. The ordering of acquisition times is indifferent, i.e., either $t_2 < t_1$ or $t_2 > t_1$ are possible cases and the order does not impact the applicability the proposed method. The problem addressed in this paper consists in detecting significant changes between these two observed images. This is a challenging task mainly due to the possible dissimilarities in terms of spatial and/or spectral resolutions and of modality. Indeed, resolution dissimilarity prevents any use of classical CD algorithms without homogenization of the resolutions as a preprocessing step (Singh, 1989; Bovolo and Bruzzone, 2015). Moreover modality dissimilarity, which makes most of the CD algorithms inoperative because their inability of handling images of different nature (Ferraris et al., 2017a,b). To alleviate this issue, this work proposes to improve and generalize the CD methods introduced by Seichepine et al. (2014), Gong et al. (2016), Lu et al. (2017). Following the widely admitted forward model described in Section 2.1 and adopting consistent notations, the observed images \mathbf{Y}_1 and \mathbf{Y}_2 can be related to two latent images $\mathbf{X}_1 \in \mathbb{R}^{L_1 \times N_1}$ and $\mathbf{X}_2 \in \mathbb{R}^{L_2 \times N_2}$

$$\mathbf{Y}_1 = T_1[\mathbf{X}_1] \quad (12a)$$

$$\mathbf{Y}_2 = T_2[\mathbf{X}_2] \quad (12b)$$

where T_1 and T_2 denote two degradation operators imposed by the sensors S_1 and S_2 . Note that (12) is a double instance of the model (1). In particular, in the CD context considered in this work, the two

latent images \mathbf{X}_1 and \mathbf{X}_2 are supposed to represent the same geographical region provided the observed images have been beforehand co-registered.

Both latent images can be represented thanks to a dedicated dictionary-based decomposition, as stated in Section 2.2. More precisely, a pair of homologous patches extracted from each image represents the same geographical spot. Each patch can be reconstructed from a sparse linear combination of atoms of an image-dependent dictionary. In the absence of changes between the two observed images, the sparse codes associated with the corresponding latent image are expected to be approximately the same and the two learned dictionaries are coupled (Yang et al., 2010, 2012; Zeyde et al., 2010). This coupling can be understood as the ability of deriving a joint representation for homologous multiple observations in a latent coupled space (Gong et al., 2016). Akin to the manifold proposed by Prendes et al. (2015), this representation offers the opportunity to analyze images of different modalities in a common dual space. In the case where a pair of homologous patches has been extracted from two images representing the same scene, given perfect estimated coupled dictionaries, each patch should be exactly reconstructed thanks to the same sparse code. In other words, the pair of patches is an element of the latent coupled space. Nevertheless, in the case where the pair of homologous patches does not represent exactly the same scene, owing to a change that occurs between acquisitions, perfect reconstruction cannot be achieved using the same code. This means that the pair of patches does not belong to the coupled spaces. Using the same code for reconstruction amounts to estimate the point in the coupled spaces that best approximates the patch pair. Thereby, relaxing this constraint in some possible change locations may provide an accurate reconstruction of both images while spatially mapping change locations. In the specific context of CD addressed in this work, this finding suggests to evaluate any change between the two observed, or equivalently latent, images by comparing the corresponding codes

$$\Delta \mathbf{A} = \mathbf{A}_2 - \mathbf{A}_1 \quad (13)$$

where $\Delta \mathbf{A} = [\Delta \mathbf{a}_1, \dots, \Delta \mathbf{a}_{N_p}]$ and $\Delta \mathbf{a}_i \in \mathbb{R}^{N_d}$ denotes the code change vector associated with the i th patch, $i = 1, \dots, N_p$. Then, to spatially locate the changes, a natural approach consists in monitoring the magnitude of $\Delta \mathbf{A}$, summarized by the code change energy image (Bovolo and Bruzzone, 2007)

$$\mathbf{e} = [e_1, \dots, e_{N_p}] \in \mathbb{R}^{N_p} \quad (14)$$

with

$$e_i = \|\Delta \mathbf{a}_i\|_2, \quad i = 1, \dots, N_p.$$

Note that, in the case of analyzing a pair of optical images, Zanetti et al. (2015) proposed to describe the components e_i of the energy vector \mathbf{e} thanks to a Rayleigh–Rice mixture model whose parameters can be estimated to locate the changes. Conversely, in this work we propose to derive the CD rule directly from this magnitude. When the CD problem in the i th patch is formulated as the binary hypothesis testing

$$\begin{cases} \mathcal{H}_{0,i} & : \text{no change occurs in the } i\text{th patch} \\ \mathcal{H}_{1,i} & : \text{a change occurs in the } i\text{th patch} \end{cases}$$

a patch-wise statistical test can be written by thresholding the code change energy

$$\begin{matrix} \mathcal{H}_{1,i} \\ e_i \geq \tau \\ \mathcal{H}_{0,i} \end{matrix}$$

where the threshold $\tau \in [0, \infty]$ implicitly adjusts the target probability of false alarm or, reciprocally, the probability of detection. The final binary CD map denoted $\mathbf{m} = [m_1, \dots, m_{N_p}] \in \{0, 1\}^{N_p}$ can be derived as

$$m_i = \begin{cases} 1 & \text{if } e_i \geq \tau \quad (\mathcal{H}_{1,i}) \\ 0 & \text{otherwise} \quad (\mathcal{H}_{0,i}). \end{cases}$$

The spatial resolution of this CD map is defined by the number N_p of homologous patches extracted from the latent images \mathbf{X}_1 and \mathbf{X}_2 . This number can be tailored by the user according to the adopted strategy of patch extraction. In practice, to reach the highest resolution, overlapping patches should be extracted according to the regular grid defined by the observed image of highest resolution, i.e., $N_p = \max\{N_1, N_2\}$.

Finally, to solve the multimodal image CD problem, the key issue lies in the joint estimation of the pair of representation codes $\{\mathbf{A}_1, \mathbf{A}_2\}$ or, equivalently, to the joint estimation of one code matrix and of the change code matrix, i.e. of $\{\mathbf{A}_1, \Delta\mathbf{A}\}$, as well as of the pair of coupled dictionary $\{\mathbf{D}_1, \mathbf{D}_2\}$ and consequently of the pair of latent images $\{\mathbf{X}_1, \mathbf{X}_2\}$ from the joint forward model (12). The next paragraph introduces the CD-driven optimization problem to be solved.

3.2. Coupled dictionary learning for CD

The single dictionary estimation problem presented on Section 2.3 can be generalized to take into account the modeling presented in Section 3.1. Nevertheless, some previous considerations must be carefully handled in order to provide good coupling of the two dictionaries.

As the prior information about the dictionaries constrains each atom into the set \mathcal{S} of unitary energy defined by (6), an unbiased estimation of the code change vector would allow a pair of unchanged homologous patches to be reconstructed with exactly the same code, while changed patches would exhibit differences in their code. Obviously, this can only be achieved if the coupled dictionaries represent data with the same dynamics and resolutions. However, when analyzing images of different modalities and/or resolutions, this assumption can be not fulfilled. To alleviate this issue, we propose to resort to the strategy proposed by Seichepine et al. (2014), by introducing an additional diagonal scaling matrix constrained to the set $\mathcal{C} \triangleq \{\mathbf{S} \in \mathbb{R}_+^{N_{d1} \times N_{d1}} \mid \mathbf{S} = \text{diag}(s), s \geq 0\}$ where N_{d1} is the size of the dictionary \mathbf{D}_1 . This scaling matrix gathers the code energy differences originated from different modalities for each pair of coupled atoms. This is essential to ensure that the sparse codes of the two observed images are directly comparable, following (13), and then properly estimated. Therefore, considering a pair of homologous patches, their joint representation model derived from (4) can be written as

$$\begin{aligned} \mathbf{p}_{1i} &= \mathcal{R}_{1i} \mathbf{X}_1 \approx \mathbf{D}_1 \mathbf{S} \mathbf{a}_{1i} \\ \mathbf{p}_{2i} &= \mathcal{R}_{2i} \mathbf{X}_2 \approx \mathbf{D}_2 \mathbf{a}_{2i} = \mathbf{D}_2 (\mathbf{a}_{1i} + \Delta \mathbf{a}_i) \end{aligned} \quad (15)$$

where $\{\mathbf{p}_{1i}, \mathbf{p}_{2i}\}$ represents the pair of homologous patches and \mathbf{S} is the diagonal scaling matrix.

Since the codes \mathbf{A}_1 and \mathbf{A}_2 are now element-wise comparable, a natural choice to enforce coupling between them should be the equality $\mathbf{A}_1 = \mathbf{A}_2 = \mathbf{A}$. This has been a classical assumption in various coupled dictionary learning applications (Yang et al., 2010; Zeyde et al., 2010; Yang et al., 2012). Nevertheless, in a CD context, some spatial positions may not contain the same objects. To account for possible changes in some specific locations while most of the patches remain unchanged, as in Ferraris et al. (2017b), the code change energy matrix \mathbf{e} defined by (14) is expected to be sparse. As a consequence, the corresponding regularizing function is chosen as the sparsity-inducing ℓ_1 -norm of the code change energy matrix \mathbf{e} or, equivalently, as the $\ell_{2,1}$ -norm of the code change matrix

$$\phi_2(\Delta\mathbf{A}) = \|\Delta\mathbf{A}\|_{2,1} = \sum_{i=1}^{N_p} \|\Delta \mathbf{a}_i\|_2. \quad (16)$$

This regularization is a specific instance of the non-overlapping group-lasso penalization (Bach, 2011) which has been considered in various applications to promote structured sparsity (Wright et al., 2009; Févotte and Dobleigeon, 2015; Ferraris et al., 2017b).

Then, a Bayesian model extending the one derived for a single image (7) leads to the posterior distribution of the parameters of interest

$$\begin{aligned} p(\mathbf{X}_1, \mathbf{X}_2, \mathbf{D}_1, \mathbf{D}_2, \mathbf{S}, \mathbf{A}_1, \Delta\mathbf{A} | \mathbf{Y}_1, \mathbf{Y}_2) \\ \propto p(\mathbf{Y}_1 | \mathbf{X}_1) p(\mathbf{Y}_2 | \mathbf{X}_2) \\ \times p(\mathbf{X}_1 | \mathbf{D}_1, \mathbf{S}, \mathbf{A}_1) p(\mathbf{X}_2 | \mathbf{D}_2, \mathbf{A}_1, \Delta\mathbf{A}) \\ \times p(\mathbf{D}_1) p(\mathbf{D}_2) p(\mathbf{S}) p(\mathbf{A}_1) p(\Delta\mathbf{A}). \end{aligned} \quad (17)$$

By incorporating all previously defined prior distributions (or, equivalently, regularizations), the joint MAP estimator $\hat{\boldsymbol{\theta}}_{\text{MAP}} = \{\hat{\mathbf{X}}_{1,\text{MAP}}, \hat{\mathbf{X}}_{2,\text{MAP}}, \hat{\mathbf{D}}_{1,\text{MAP}}, \hat{\mathbf{D}}_{2,\text{MAP}}, \hat{\mathbf{S}}_{\text{MAP}}, \hat{\mathbf{A}}_{1,\text{MAP}}, \Delta \hat{\mathbf{A}}_{\text{MAP}}\}$ of the quantities of interest can be obtained by minimizing the negative log-posterior, leading to the following minimization problem

$$\hat{\boldsymbol{\theta}}_{\text{MAP}} \in \underset{\boldsymbol{\theta}}{\text{argmin}} \mathcal{J}(\boldsymbol{\theta}) \quad (18)$$

with

$$\begin{aligned} \mathcal{J}(\boldsymbol{\theta}) &\triangleq \mathcal{J}(\mathbf{X}_1, \mathbf{X}_2, \mathbf{D}_1, \mathbf{D}_2, \mathbf{S}, \mathbf{A}_1, \Delta\mathbf{A}) \\ &= D(\mathbf{Y}_1 | \mathbf{X}_1) + D(\mathbf{Y}_2 | \mathbf{X}_2) \\ &+ \frac{\sigma_1^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_{1i} \mathbf{X}_1 - \mathbf{D}_1 \mathbf{S} \mathbf{a}_{1i}\|_{\text{F}}^2 + \Psi(\mathbf{X}_1) \\ &+ \frac{\sigma_2^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_{2i} \mathbf{X}_2 - \mathbf{D}_2 (\mathbf{a}_{1i} + \Delta \mathbf{a}_i)\|_{\text{F}}^2 + \Psi(\mathbf{X}_2) \\ &+ \lambda \|\mathbf{A}_1\|_1 + \lambda \|\mathbf{A}_1 + \Delta\mathbf{A}\|_1 + \gamma \|\Delta\mathbf{A}\|_{2,1} \\ &+ \iota_{\mathcal{S}}(\mathbf{D}_1) + \iota_{\mathcal{S}}(\mathbf{D}_2) + \iota_{\mathcal{C}}(\mathbf{S}). \end{aligned} \quad (19)$$

The next section describes an iterative algorithm which solves the minimization problem in (18).

4. Minimization algorithm

Given the nature of the optimization problem (18), which is genuinely nonconvex and nonsmooth, the adopted minimization strategy relies on the proximal alternating linearized minimization (PALM) scheme (Bolte et al., 2014). PALM is an iterative, gradient-based algorithm which generalizes the Gauss–Seidel method. It performs iterative proximal gradient steps w.r.t. each block of variables from $\boldsymbol{\theta}$ and ensures convergence to a local critical point $\boldsymbol{\theta}^*$. It has been successfully applied in many matrix factorization cases (Bolte et al., 2014; Cavalcanti et al., 2017; Thouvenin et al., 2016). Now, the goal is to generalize the single factorization to coupled factorization. The resulting CD-driven coupled dictionary learning (CDL) algorithm, whose main steps are described in the following paragraphs, is summarized in Algorithm 1.

4.1. PALM implementation

The PALM algorithm was proposed by Bolte et al. (2014) for solving a broad class of problems involving the minimization of the sum of finite collections of possibly nonconvex and nonsmooth functions. Particularly, the target optimization function is composed by a coupling function gathering the block of variables, denoted $H(\cdot)$, and regularization functions for each block. Nonconvexity constraint is assumed for either coupling or regularization functions. One of the main advantages of the PALM algorithm over classical optimization algorithms is that each bounded sequence generated by PALM converges to a critical point. The rationale of the method can be seen as an alternating minimization approach for the proximal forward-backward algorithm (Combettes and Wajs, 2005). Some assumptions are required in order to solve this problem with all guarantees of convergence (c.f (Bolte et al., 2014, Assumption 1, Assumption 2)). The most restrictive one (Bolte et al., 2014, Assumption 2(ii)) requires that the partial gradient of the coupling function $H(\cdot)$ is globally Lipschitz continuous for each block of variable keeping the remaining ones fixed.

Algorithm 1: PALM-CDL**Data:** \mathbf{Y} **Input:** $\mathbf{A}_1^{(0)}, \Delta\mathbf{A}^{(0)}, \mathbf{D}_1^{(0)}, \mathbf{D}_2^{(0)}, \mathbf{S}^{(0)}, \mathbf{X}_1^{(0)}, \mathbf{X}_2^{(0)}$ $k \leftarrow 0$ **begin****while** *stopping criterion not satisfied* **do**

// Code update

 $\mathbf{A}^{(k+1)} \leftarrow \text{Update}(\mathbf{A}^{(k)})$ // cf. (21) $\Delta\mathbf{A}^{(k+1)} \leftarrow \text{Update}(\Delta\mathbf{A}^{(k)})$ // cf. (24)

// Dictionary update

 $\mathbf{D}_1^{(k+1)} \leftarrow \text{Update}(\mathbf{D}_1^{(k)})$ // cf. (27) $\mathbf{D}_2^{(k+1)} \leftarrow \text{Update}(\mathbf{D}_2^{(k)})$ // cf. (27)

// Scale update

 $\mathbf{S}^{(k+1)} \leftarrow \text{Update}(\mathbf{S}^{(k)})$ // cf. (30)

// Latent image update

 $\mathbf{X}_1^{(k+1)} \leftarrow \text{Update}(\mathbf{X}_1^{(k)})$ // cf. (33) $\mathbf{X}_2^{(k+1)} \leftarrow \text{Update}(\mathbf{X}_2^{(k)})$ // cf. (33) $k \leftarrow k + 1$ $\hat{\mathbf{A}}_1 \leftarrow \mathbf{A}_1^{(k+1)}, \Delta\hat{\mathbf{A}} \leftarrow \Delta\mathbf{A}^{(k+1)},$ $\hat{\mathbf{D}}_1 \leftarrow \mathbf{D}_1^{(k+1)}, \hat{\mathbf{D}}_2 \leftarrow \mathbf{D}_2^{(k+1)},$ $\hat{\mathbf{S}} \leftarrow \mathbf{S}^{(k+1)},$ $\hat{\mathbf{X}}_1 \leftarrow \mathbf{X}_1^{(k+1)}, \hat{\mathbf{X}}_2 \leftarrow \mathbf{X}_2^{(k+1)}$ **Result:** $\hat{\mathbf{A}}_1, \Delta\hat{\mathbf{A}}, \hat{\mathbf{D}}_1, \hat{\mathbf{D}}_2, \hat{\mathbf{S}}, \hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2$

Indeed, it is a classical assumption for proximal gradient methods which guarantees a sufficient descent property.

Therefore, given the objective function to be minimized (19) and considering the same structure proposed by Bolte et al. (2014) and the Lipschitz property for linear combinations of functions (Eriksson et al., 2004), let us define the coupling function $H(\Theta)$ as

$$\begin{aligned} H(\Theta) &\triangleq H(\mathbf{X}_1, \mathbf{X}_2, \mathbf{D}_1, \mathbf{D}_2, \mathbf{S}, \mathbf{A}_1, \Delta\mathbf{A}) \\ &= \Psi(\mathbf{X}_1) + \Psi(\mathbf{X}_2) + \frac{\sigma_1^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_{1i}\mathbf{X}_1 - \mathbf{D}_1\mathbf{S}\mathbf{a}_{1i}\|_{\mathbb{F}}^2 \\ &\quad + \frac{\sigma_2^2}{2} \sum_{i=1}^{N_p} \|\mathcal{R}_{2i}\mathbf{X}_2 - \mathbf{D}_2(\mathbf{a}_{1i} + \Delta\mathbf{a}_i)\|_{\mathbb{F}}^2 + \lambda \|\mathbf{A}_1 + \Delta\mathbf{A}\|_1. \end{aligned} \quad (20)$$

This coupling function defined accordingly does not fulfill (Bolte et al., 2014, Assumption 2(ii)) because some of its terms are nonsmooth, specifically the TV regularizations $\Psi(\cdot)$ and the ℓ_1 -norm sparsity promoting regularizations applied to \mathbf{A}_2 . Thus, to ensure such a coupling function is in agreement with the required assumptions, smooth relaxations of $\Psi(\cdot)$ and $\|\cdot\|_1$ are applied by using the pseudo-Huber function (Fountoulakis and Gondzio, 2016; Jensen et al., 2012).

The remaining terms of (19) are composed of the regularization functions associated with each variable block. Within the PALM structure, a gradient step applied to the coupling function w.r.t. a given variable block is followed by proximal step associated with the corresponding regularization functions. As a consequence, those regularization functions must be proximal-like where their proximal mappings or projections must exist and have closed-form solutions. It is important to keep in mind that, even if the convergence is guaranteed for all optimization orderings, it should not vary during iterations. Thus, the updating rules for each optimization variable in Algorithm 1 are defined. More details about the proximal operators and projections involved in this section are given in Appendix B.

4.2. Optimization with respect to \mathbf{A}_1

Considering the single block optimization variable \mathbf{A}_1 , and assuming that the remaining variables are fixed, the PALM updating step can be

written

$$\mathbf{A}_1^{(k+1)} = \text{prox}_{\lambda\|\cdot\|_1 \geq 0}^{L_{\mathbf{A}_1}} \left(\mathbf{A}_1^{(k)} - \frac{1}{L_{\mathbf{A}_1}^{(k)}} \nabla_{\mathbf{A}_1} H(\Theta) \right) \quad (21)$$

with

$$\begin{aligned} \nabla_{\mathbf{A}_1} H(\Theta) &= \sigma_1^2 \mathbf{S}^T \mathbf{D}_1^T (\mathbf{D}_1 \mathbf{S} \mathbf{A}_1 - \mathbf{P}_1) \\ &\quad + \sigma_2^2 \mathbf{D}_2^T (\mathbf{D}_2 (\mathbf{A}_1 + \Delta\mathbf{A}) - \mathbf{P}_2) \\ &\quad + \lambda \frac{[\mathbf{A}_1 + \Delta\mathbf{A}]_i}{\sqrt{[\mathbf{A}_1 + \Delta\mathbf{A}]_i^2 + \epsilon_{\mathbf{A}_1}^2}} \end{aligned} \quad (22)$$

where $[\cdot]_i / [\cdot]_i$ should be understood as an element-wise operation and $L_{\mathbf{A}_1}^{(k)}$ is the associated Lipschitz constant

$$L_{\mathbf{A}_1}^{(k)} = \sigma_1^2 \|\mathbf{S}^T \mathbf{D}_1^T \mathbf{D}_1 \mathbf{S}\| + \sigma_2^2 \|\mathbf{D}_2^T \mathbf{D}_2\| + \frac{\lambda}{\epsilon_{\mathbf{A}_1}}. \quad (23)$$

Note that $\text{prox}_{\lambda\|\cdot\|_1 \geq 0}^{L_{\mathbf{A}_1}}(\cdot)$ can be simply computed by considering the positive part of the soft-thresholding operator (Parikh et al., 2014).

4.3. Optimization with respect to $\Delta\mathbf{A}$

Similarly, considering the single block optimization variable $\Delta\mathbf{A}$ and consistent notations, the PALM update can be derived as

$$\Delta\mathbf{A}^{(k+1)} = \text{prox}_{\|\cdot\|_{2,1}}^{L_{\Delta\mathbf{A}}^{(k)}} \left(\Delta\mathbf{A}^{(k)} - \frac{1}{L_{\Delta\mathbf{A}}^{(k)}} \nabla_{\Delta\mathbf{A}} H(\Theta) \right) \quad (24)$$

where

$$\begin{aligned} \nabla_{\Delta\mathbf{A}} H(\Theta) &= \sigma_2^2 \mathbf{D}_2^T (\mathbf{D}_2 (\mathbf{A}_1 + \Delta\mathbf{A}) - \mathbf{P}_2) \\ &\quad + \lambda \frac{[\mathbf{A}_1 + \Delta\mathbf{A}]_i}{\sqrt{[\mathbf{A}_1 + \Delta\mathbf{A}]_i^2 + \epsilon_{\mathbf{A}_1}^2}} \end{aligned} \quad (25)$$

and

$$L_{\Delta\mathbf{A}}^{(k)} = \sigma_2^2 \|\mathbf{D}_2^T \mathbf{D}_2\| + \frac{\lambda}{\epsilon_{\mathbf{A}_1}}. \quad (26)$$

The proximal operator $\text{prox}_{\|\cdot\|_{2,1}}^{L_{\Delta\mathbf{A}}^{(k)}}(\cdot)$ can be simply computed as a group soft-thresholding operator (Ferraris et al., 2017b), where each group is composed by each column of $\Delta\mathbf{A}$.

4.4. Optimization with respect to \mathbf{D}_α

As before, considering the single block optimization variable \mathbf{D}_α with $\alpha = \{1, 2\}$, the PALM updating steps can be written as

$$\mathbf{D}_\alpha^{(k+1)} = \mathcal{P}_{\mathcal{S}} \left(\mathbf{D}_\alpha^{(k)} - \frac{1}{L_{\mathbf{D}_\alpha}^{(k)}} \nabla_{\mathbf{D}_\alpha} H(\Theta) \right) \quad (27)$$

where

$$\nabla_{\mathbf{D}_\alpha} H(\Theta) = \sigma_\alpha^2 (\mathbf{D}_\alpha \bar{\mathbf{A}}_\alpha - \mathbf{P}_\alpha) \bar{\mathbf{A}}_\alpha^T \quad (28)$$

and $L_{\mathbf{D}_\alpha}^{(k)}$ is the Lipschitz constant

$$L_{\mathbf{D}_\alpha}^{(k)} = \sigma_\alpha^2 \|\bar{\mathbf{A}}_\alpha \bar{\mathbf{A}}_\alpha^T\| \quad (29)$$

with $\bar{\mathbf{A}}_1 = \mathbf{S}\mathbf{A}_1$ and $\bar{\mathbf{A}}_2 = \mathbf{A}_1 + \Delta\mathbf{A}$. Note that the projection $\mathcal{P}_{\mathcal{S}}(\cdot)$ can be computed as in Mairal et al. (2009), keeping only the values greater than zero.

4.5. Optimization with respect to \mathbf{S}

The updating rule of the scaling matrix \mathbf{S} can be written as

$$\mathbf{S}^{(k+1)} = \mathcal{P}_{\mathcal{C}} \left(\mathbf{S}^{(k)} - \frac{1}{L_{\mathbf{S}^{(k)}}} \nabla_{\mathbf{S}} H(\Theta) \right) \quad (30)$$

where

$$\nabla_{\mathbf{S}} H(\boldsymbol{\Theta}) = \sigma_1^2 \mathbf{D}_1^T (\mathbf{D}_1 \mathbf{S} \mathbf{A}_1 - \mathbf{P}_1) \mathbf{A}_1^T \quad (31)$$

and $L_{\mathbf{S}}^{(k)}$ is the Lipschitz constant related to $\nabla_{\mathbf{S}} f(\boldsymbol{\Theta})$

$$L_{\mathbf{S}}^{(k)} = \sigma_1^2 \left\| \mathbf{D}_1^T \mathbf{D}_1 \mathbf{A}_1 \mathbf{A}_1^T \right\|. \quad (32)$$

The projection $\mathcal{P}_C(\cdot)$ constrains all diagonal elements of \mathbf{S} to be nonzero.

4.6. Optimization with respect to \mathbf{X}_α

Finally, the updates of the latent images \mathbf{X}_α ($\alpha \in \{1, 2\}$) are achieved as follows

$$\mathbf{X}_\alpha^{(k+1)} = \text{prox}_{D_\alpha(\mathbf{Y}_\alpha|\cdot)}^{L_{\mathbf{X}_\alpha}^{(k)}} \left(\mathbf{X}_\alpha^{(k)} - \frac{1}{L_{\mathbf{X}_\alpha}^{(k)}} \nabla_{\mathbf{X}_\alpha} H(\boldsymbol{\Theta}) \right) \quad (33)$$

with

$$\begin{aligned} \nabla_{\mathbf{X}_\alpha} H(\boldsymbol{\Theta}) &= \sigma_\alpha^2 \sum_{i=1}^{N_p} \mathcal{R}_{ai}^T (\mathcal{R}_{ai} \mathbf{X}_\alpha - \mathbf{D}_\alpha \bar{\mathbf{a}}_{ai}) \\ &\quad - \tau_\alpha \text{div} \left(\frac{[\nabla \mathbf{X}_\alpha]_i}{\sqrt{[\nabla \mathbf{X}_\alpha]_i^2 + \epsilon_{\mathbf{X}_\alpha}^2}} \right) \end{aligned} \quad (34)$$

and

$$L_{\mathbf{X}_\alpha}^{(k)} = \sigma_\alpha^2 \left\| \sum_{i=1}^{N_p} \mathcal{R}_{ai}^T \mathcal{R}_{ai} \right\| + \frac{8\tau_\alpha}{\epsilon_{\mathbf{X}_\alpha}^2} \quad (35)$$

and where $\text{div}(\cdot)$ stands for the discrete divergence (Chambolle, 2004).

Note that, $\text{prox}_{D_\alpha(\mathbf{Y}_\alpha|\cdot)}^{L_{\mathbf{X}_\alpha}^{(k)}}$ represents the proximal mapping for the divergence measure associated with the likelihood function characterizing the modality of the observed image \mathbf{Y}_α . For the most common remote sensing modalities, e.g., optical and radar, these divergences are well documented and Appendix A presents the corresponding proximal operators.

5. Performance analysis

Real datasets with corresponding ground truth are too scarce to statistically assess the performance of CD algorithms. Indeed, this assessment would require a huge number of pairs of images acquired at two different dates, geometrically co-registered and presenting changes. These pairs should also be accompanied by a ground truth (i.e., a binary CD mask locating the actual changes) to allow quantitative figures-of-merit to be computed. As a consequence, in a first step, we illustrate the algorithm and state-of-the-art method outcomes over such rare examples corresponding to real images, real changes and associated ground truth (Section 5.3.1). This first set of experiments is conducted on images of same spatial resolutions. Thus we also exhibit another set of examples that involves images with different resolutions, but alas without ground truth. For this second example, the accuracy of the proposed method cannot be quantified, but can be evaluated through visual inspection (Section 5.3.2). In a second step, to conduct a complete performance analysis, the algorithm and comparable methods will be tested on image pairs that are representative of possible scenarios considered in this paper, i.e., coming from multimodal images. This test set is composed of real images, however simulated changes and associated ground truth (Section 5.4).

5.1. Compared methods

As the number of unsupervised multimodal CD methods is rather reduced, the proposed technique has been compared to the unsupervised fuzzy-based (F) method proposed by Gong et al. (2016), that is able to deal with multimodal images, to the robust fusion (RF)

method proposed by Ferraris et al. (2017b) which deals exclusively with multi-band optical images and with unsupervised segmentation-based (S) technique proposed by Huang et al. (2015). The fuzzy-based method by Gong et al. (2016) relies on a coupled dictionary learning methodology using a modified K-SVD (Aharon et al., 2006) with an iterative patch selection procedure to provide only unchanged patches for the coupled dictionary training phase. Then, the sparse code for each observed image is estimated separately from each other allowing to compute the cross-image reconstruction errors. Finally, a local fuzzy C-Means is applied to the mean of the cross-image reconstruction errors in order to separate change and unchanged classes. Equivalently to the proposed one, this method makes no assumption about the joint observation model. On the other hand, the robust fusion method by Ferraris et al. (2017b) is based on a more constrained joint observation model, considering that the two latent images share the same resolutions and differ only in changed pixels. Finally, the method proposed by Huang et al. (2015) replaces the pixel-based approach used on all previous methods to a feature-based approach. In this approach, features are derived from the segmentation of each image. Then, a difference map is generated based on metrics computed for the matching features extracted on previous steps at different scales. This strategy is used in order to provide finer details. At the end, the change map is generated using the histogram of difference map. The final change maps estimated by these algorithms are denoted as $\hat{\mathbf{m}}_F$, $\hat{\mathbf{m}}_{RF}$ and $\hat{\mathbf{m}}_S$, respectively, while the proposed PALM-CDL method provides a change map denoted $\hat{\mathbf{m}}_{CDL}$.

5.2. Figures-of-merit

The CD performance of the different methods has been assessed through the empirical receiver operating characteristics (ROC) curves, representing the estimated pixel-wise probability of detection (PD) as a function of the probability of false alarm (PFA). Moreover, two quantitative criteria derived from these ROC curves have been computed, namely, (i) the area under the curve (AUC), corresponding to the integral of the ROC curve and (ii) the distance between the no detection point (PFA = 1, PD = 0) and the point at the interception of the ROC curve with the diagonal line defined by PFA = 1 - PD. For both metrics, the greater the criterion, the better the detection.

5.3. Illustration through real images with real changes

In a first step, experiments are conducted on real images with real changes to emphasize the reliability of the proposed CD method and to illustrate the performance of the proposed algorithmic framework. Three distinct scenarios involving 3 pairs of images of different modalities and resolutions, are considered namely,

- Scenario 1 considers two optical images: the acquisition process is very similar for the two images and the image formation processes is characterized by an additive Gaussian noise corruption for both sensors.
- Scenario 2 considers two SAR images: the image formation process is not the same as for optical images, in particular differing on the noise model, i.e., multiplicative Gamma noise instead of additive Gaussian model.
- Scenario 3 considers a SAR image and an optical image: for this more challenging situation, there is no similarity between the noise corruption models for the two sensors.

To summarize, Scenarios 1 and 2 are dedicated to a pair of images with the same modality, but with a variation on the properties of images between scenarios, e.g., the noise statistics. Note that the proposed CDL algorithm has not been designed to specifically handle these conventional scenarios. However, they are still considered to evaluate the performance of the proposed method, in particular w.r.t. the methods specifically designed to address scenario 1 or 2. Conversely, Scenario 3 handles images of different modalities. All considered images have been manually geographically aligned to fulfill the requirements imposed by the considered CD setup.

Table 1

Real images affected by real changes with ground truth for Scenarios 1–3: quantitative detection performance (AUC and distance).

		\hat{m}_F	\hat{m}_{RF}	\hat{m}_S	\hat{m}_{CDL}
Sc. 1	AUC	0.870426	0.866061	0.750601	0.987379
	Dist.	0.831983	0.788579	0.692469	0.950695
Sc. 2	AUC	0.823414	0.93982	0.874743	0.981355
	Dist.	0.757076	0.869387	0.80188	0.954995
Sc. 3	AUC	0.818246	0.862729	0.862025	0.966283
	Dist.	0.769877	0.79658	0.80078	0.912191

5.3.1. Case of same resolutions and ground truth

For this first set of experiments, images of same spatial resolutions are considered. They are accompanied by a ground truth in the form of an actual change map m to be estimated.

Scenario 1: optical vs. optical — The observed images are two multispectral (MS) optical images with 3 channels representing an urban region in the south of Toulouse, France, before (Fig. 1(a)) and after (Fig. 1(b)) the construction of a road. These 960×1560 -pixel images are both characterized by a 50cm spatial resolution. The ground-truth change mask m is represented in Fig. 1(c). Fig. 1 depicts the observed images at each date, the ground-truth change mask and the change maps estimated by the four compared methods.

The quantitative results for Scenario 1 are reported in Table 1 (lines 1 and 2) and the corresponding ROC curves in Fig. 2(a). The analysis of these results shows that the proposed method outperforms state-of-the-art methods for this scenario which involves common changes in urban areas and in MS optical images. Note that, besides the changes, this kind of situation involves a lot of small differences between the two observed images due to the variations in sun the illumination, in the vegetation cover, etc. These effects sometimes are classified as changes, increasing the false alarm rate, especially for state-of-the-art methods. Besides, the proposed method still provides the best detection for this dataset.

Scenario 2: SAR vs. SAR — For this scenario, two intensity radar images acquired over the Lake Mulargia region in Sardegna, by the Sentinel-1 satellite in 05/21/2016 (Fig. 3(a)) and 10/30/2016 (Fig. 3(b)) are considered. These 1200×1800 -pixel images are both characterized by a 10m spatial resolution. This dataset mostly presents seasonal changes, in particular the variation of flooding areas around the Lake Mulargia. The ground-truth change mask is represented in Fig. 3(c). Fig. 3 depicts the two observed images, the ground-truth change mask and the change maps estimated by the four compared methods.

The quantitative results for Scenario 2 are reported in Table 1 (lines 3 and 4) and the corresponding ROC curves in Fig. 2(b). The analysis of these results shows that the proposed method also outperforms the state-of-the-art methods for this scenario. It is a good indication of its flexibility w.r.t. image modalities. Note that, due to the multiplicative noise and the consequent strong fluctuations, the state-of-the-art methods present a lot of false alarms. This effect seems to be attenuated by the proposed method, probably thanks to the TV regularization.

Scenario 3: optical vs. SAR — In order to test the performance of the different compared methods in a multi-modality situation, we consider two images acquired over the Gloucester region, UK, before and after a catastrophic flooding accident in 2007. The before-flooding image, presented in Fig. 4(a), is a multispectral optical image with 3 channels acquired by Google Earth while the after-flooding image, depicted in Fig. 4(b), is a radar image acquired by TerraSAR-X. These 2325×4133 -pixel images are both characterized by a 7.3 m spatial resolution. The ground-truth change mask is represented on Fig. 4(c). Fig. 4 depicts the observed images at each date, the ground-truth change mask and the change maps estimated by the four comparative methods.

Table 1 (lines 5 and 6) reports the quantitative results for Scenario 3 and the corresponding ROC curves are displayed in Fig. 2(c). Similarly

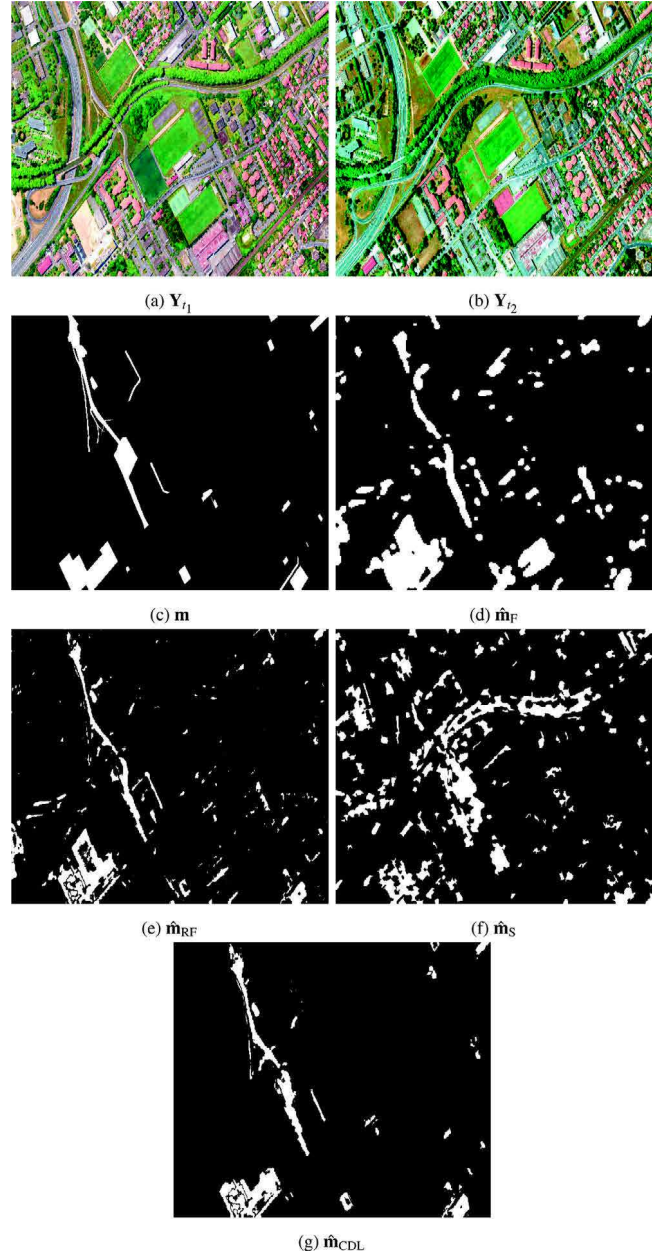


Fig. 1. Real images affected by real changes with ground truth, Scenario 1: (a) observed MS optical image Y_{t_1} from the south of Toulouse acquired before the construction of a new road, (b) observed MS optical image Y_{t_2} acquired after its construction, (c) ground-truth mask m indicating changed areas constructed by photointerpretation, (d) change map \hat{m}_F of the fuzzy method, (e) change map \hat{m}_{RF} of the robust fusion method, (f) change map \hat{m}_S of the segmentation-based method, and (g) change map \hat{m}_{CDL} of proposed method.

as in the previous scenarios, the analysis of these results shows that the proposed method outperforms the state-of-the-art methods even in this more challenging situation involving different image modalities with changes in rural and urban areas. As for Scenario 2, the TV regularization seems to be beneficial to smooth the fluctuations due to the nature of the noise in radar images, which may affect the other compared methods producing more false alarms.

Through this first illustration, we can observe that the segmentation-based method severely underperforms the other methods for scenarios 1 and 2 and underperforms the proposed method under scenario 1. For brevity, in the following, we will compare only the F and RF methods to the proposed CDL one.

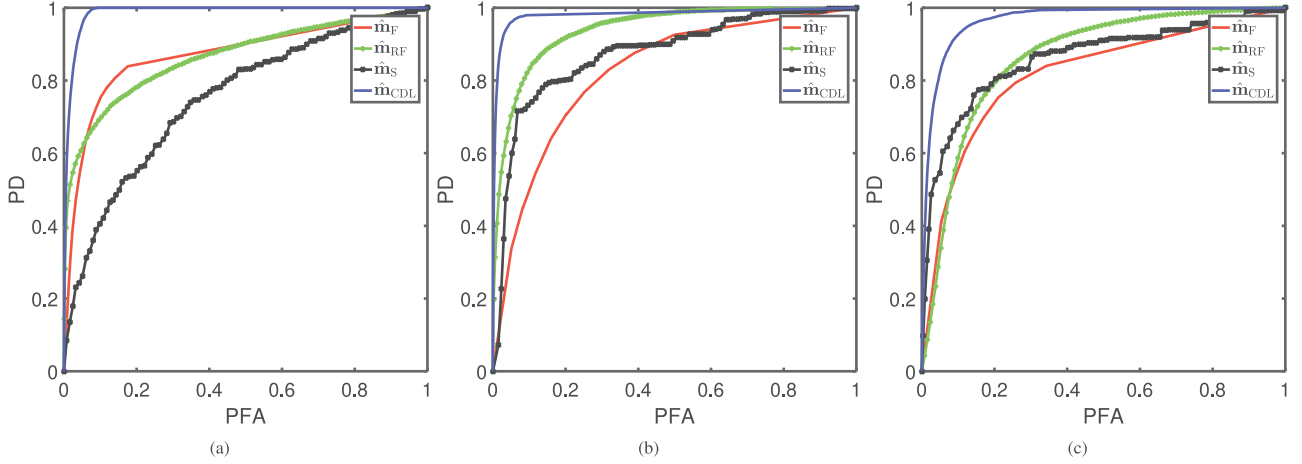


Fig. 2. Real images affected by real changes with ground truth: ROC curves for (a) Scenario 1, (b) Scenario 2, (c) Scenario 3.

5.3.2. Case of different resolutions without ground truth

The previous set of experiments considered pair of images characterized by the same spatial resolution. As a complementary analysis, this section reports experiments conducted on real images of different spatial resolutions with real changes. However, for these 3 pairs of images, corresponding to the three scenarios, no ground truth is available. We first consider a Sentinel-1 SAR image (European Space Agency, 2017a) acquired on October 28th 2016. This image is a 540×525 interferometric wide swath high resolution ground range detected multi-looked SAR intensity image with a spatial resolution of 10m according to 5 looks in the range direction. Moreover, we also consider two multispectral Landsat 8 (United States Geological Survey, 2017) 180×175 -pixel images with 30m spatial resolution and composed of the RGB visible bands (Band 2 to 4), acquired over the same region on April 15th 2015 and on September 22th 2015, respectively. Unfortunately, no ground-truth information is available for the chosen dates, as experienced in numerous experimental situations (Bovolo and Bruzzone, 2015). However, this region is characterized by interesting natural meteorological changes occurring along the seasons (e.g., drought of the Mud Lake, snow falls and vegetation growth), which helps to visually infer the major changes between observed images and to assess the relevance of the detected changes. All considered images have been manually geographically and geometrically aligned to fulfill the requirements imposed by the considered CD setup. Each scenario is individually studied considering the same denominations as in Section 5.3 and the same compared methods as in Section 5.1.

Scenario 1: optical vs. optical — In this scenario, two different situations are going to be explored, namely, observed images with the same or different resolutions. The first case considers both Landsat 8 images. Fig. 5 depicts the two observed images and the change maps estimated by the three compared methods. These change maps have been generated according to (12) where the threshold has been adjusted such that each method reveals the most important changes, i.e., the drought of the Mud Lake. As expected, the robust fusion method presents better accuracy in detection since it was specifically designed to handle such a scenario. Nevertheless, the proposed method exhibits very similar results. It is worth noting that some of the observed differences are due to the patch decomposition required by the proposed method. The fuzzy method is able to localize the strongest changes, but low energy changes are not detected. The fuzzy method also suffers from resolution loss due to the size of the patches. Contrary to the proposed method, it does not take the patch overlapping into account, which contributes to decrease the detection accuracy.

Under the same scenario (i.e. optical vs. optical), an additional pair of observed images is used to better understand the algorithm behavior when facing to images of the same modality but with different spatial

resolutions. The observed image pair is composed of the Sentinel-2 image acquired on April 12th 2016 and the Landsat 8 image acquired in September 22th 2015. Note that the two observed images have the same spectral resolution, but different spatial resolutions. Fig. 6 depicts the observed images as well as the change maps estimated by the comparative methods. Once again, it is possible to state the similarity of the results provided by the robust fusion method and the proposed one. It also shows the very poor detection performance of the fuzzy method. This may be explained by the difficulty of coupling due to differences in resolutions.

Scenario 2: SAR vs. SAR — In this scenario, observed SAR images acquired by the same sensor (Sentinel-1) are used to assess the performance of the fuzzy method and the proposed one. The robust fusion method has not been considered due to the poor results obtained on the synthetic dataset (see Section 5.4 below). Fig. 7 presents the observed images at each date and the change maps recovered by the two compared methods. The same strategy of threshold selection as for Scenario 1 has been adopted to reveal the most important changes. As expected, the proposed method presents a higher accuracy in detection than the fuzzy method. Possible reasons that may explain this difference are (i) the fuzzy method is unable to handle overlapping patches and (ii) the fuzzy method does not exploit appropriate data-fitting terms, in opposite to the proposed one. Besides, as SAR images present strong fluctuations due to their inherent image formation process, the additional TV regularization of the proposed method may contribute to smooth such fluctuations and better couple the dictionaries.

Scenario 3: optical vs. SAR — For this scenario, once again, two different situations are addressed: images with the same or different spatial resolutions. The first one considers the Sentinel-2 MS image acquired on April 12th 2016 and the Sentinel-1 SAR image acquired in October 28th 2016. Fig. 8 presents the observed images and the change maps derived from the fuzzy and proposed methods. To derive the change maps, the thresholding strategy is the same as for all previous scenarios. Once again, the proposed method shows better detection accuracy performance than the fuzzy one. It is important to emphasize the similarity of the results achieved in Scenario 3 and Scenario 2 for images acquired at the same date. Note also that this similarity can be observed for the proposed method, which contributes to increase its reliability for CD between multimodal images.

The second observed image pair consists in a Sentinel-1 SAR image acquired on April 12th 2016 and a Landsat 8 MS image acquired on September 22th 2015. This pair represents the most challenging situation among all presented images, namely differences in both modalities and resolutions. Fig. 9 presents the observed images at each date and the recovered change maps. For this last experiment, the proposed method presents better accuracy in detection than the fuzzy one. All

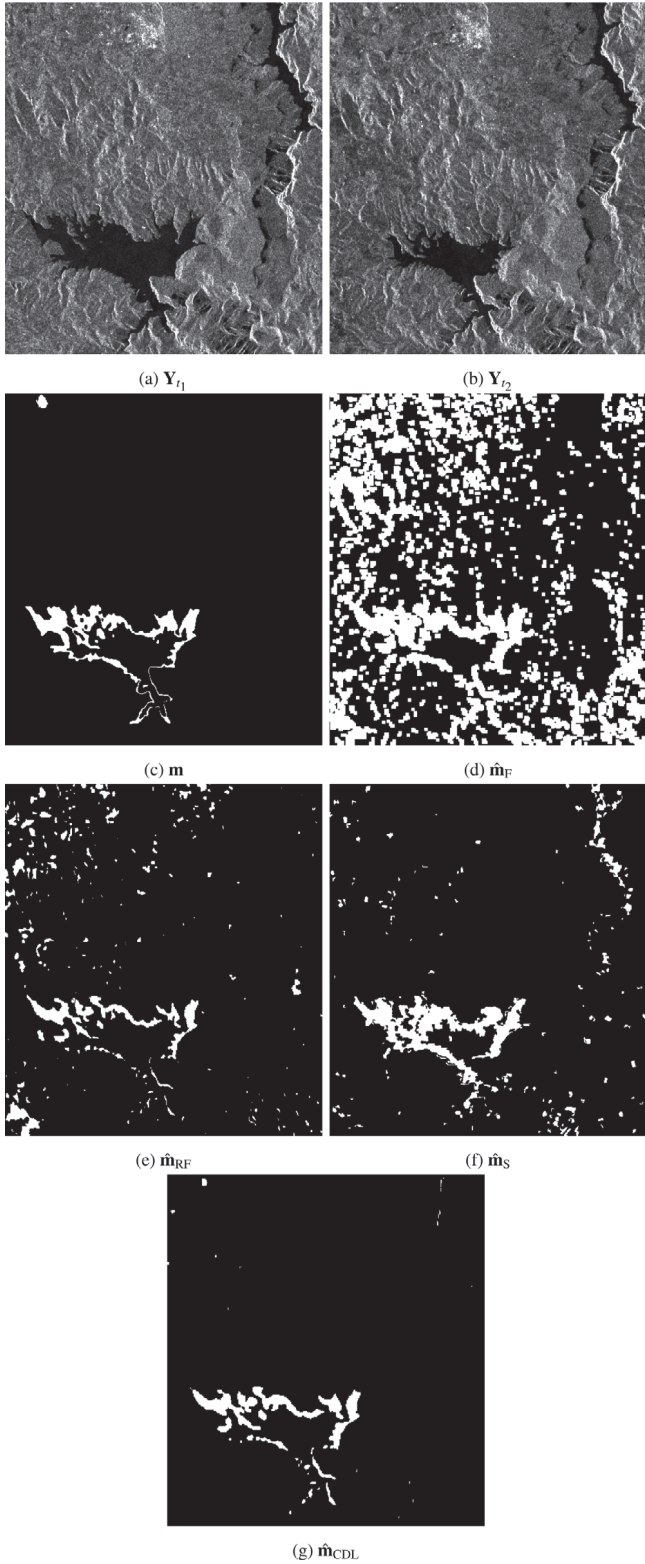


Fig. 3. Real images affected by real changes with ground truth, Scenario 2: (a) observed radar image Y_{t_1} from the Lake Mulargia acquired in 05/21/2016 by Sentinel 1, (b) observed radar image Y_{t_2} from the Lake Mulargia acquired in 10/30/2016 by Sentinel 1, (c) ground-truth mask \mathbf{m} indicating changed areas constructed by photointerpretation, (d) change map $\hat{\mathbf{m}}_F$ of the fuzzy method, (e) change map $\hat{\mathbf{m}}_{RF}$ of the robust fusion method, (f) change map $\hat{\mathbf{m}}_S$ of the segmentation-based method, and (g) change map $\hat{\mathbf{m}}_{CDL}$ of proposed method.

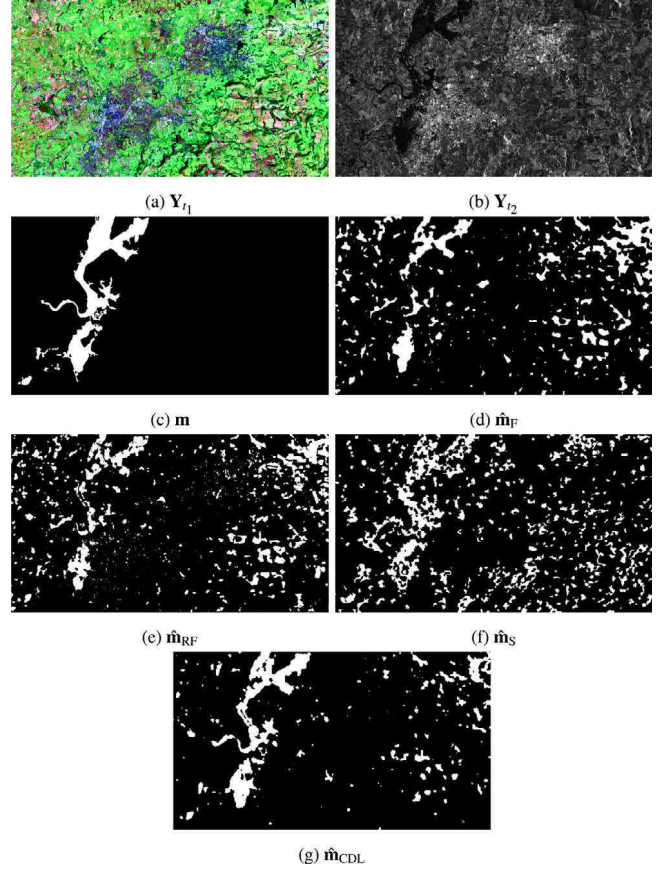


Fig. 4. Real images affected by real changes with ground truth, Scenario 3: (a) observed MS optical image Y_{t_1} from Gloucester region acquired before the flooding by Google Earth, (b) observed radar image Y_{t_2} from Gloucester region acquired after the flooding by TerraSAR-X, (c) ground-truth mask \mathbf{m} indicating changed areas constructed by photointerpretation, (d) change map $\hat{\mathbf{m}}_F$ of the fuzzy method, (e) change map $\hat{\mathbf{m}}_{RF}$ of the robust fusion method, (f) change map $\hat{\mathbf{m}}_S$ of the segmentation-based method, and (g) change map $\hat{\mathbf{m}}_{CDL}$ of proposed method.

differences in all previous situations can be observed in this scenario, culminating in the poor detection performance of the fuzzy method and a reliable change map for the proposed one.

5.4. Statistical performance assessment

Finally, the last set of experiments aims at statistically evaluating the performance of compared algorithms thanks to simulations on real images affected by synthetic changes. More precisely, in the case of multi-band images, a dedicated CD evaluation protocol was proposed by Ferraris et al. (2017a) based on a single high spatial resolution hyperspectral reference image. The experiments conducted in this work follow the same strategy. Two multimodal reference images acquired at the same date have been selected as change-free latent images. By conducting simple copy-paste of regions, as in Ferraris et al. (2017a), changes have been generated in both images as well as their corresponding ground-truth maps. This process allows synthetic yet realistic changes to be incorporated within one of these latent images, w.r.t. a pre-defined binary reference change mask locating the pixels affected by these changes and further used to assess the performance of the CD algorithms. This process is detailed in what follows.

5.4.1. Simulation protocol

Reference images — The reference images X_1^{ref} and X_2^{ref} used in this experiment comes from two largely studied open access satellite sensors, namely Sentinel-1 (European Space Agency, 2017a) and

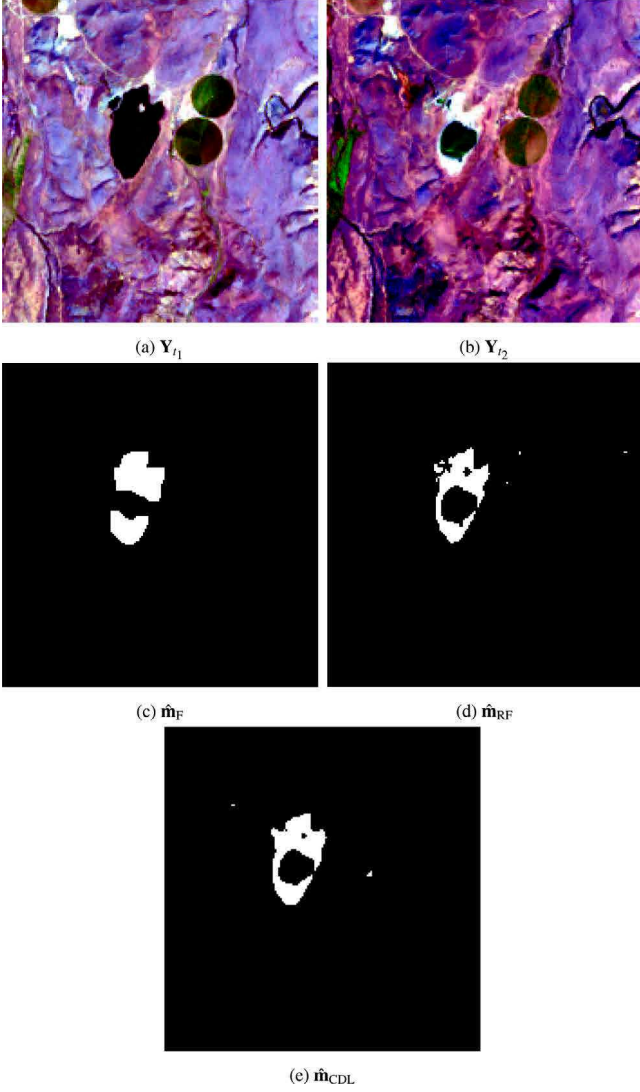


Fig. 5. Real images affected by real changes without ground truth, Scenario 1 (same spatial resolutions): (a) observed Landsat 8 MS image Y_{t_1} acquired on 04/15/2015, (b) Landsat 8 MS image Y_{t_2} acquired on 09/22/2015, (c) change map \hat{m}_F of the fuzzy method, (d) change map \hat{m}_{RF} of the robust fusion method and (e) change map \hat{m}_{CDL} of the proposed method.

Sentinel-2 (European Space Agency, 2017b) operated by the European Spatial Agency. These images have been acquired over the same geographical area, i.e., the Mud Lake region in Lake Tahoe, at the same date on April 12th 2016. To fulfill the requirements imposed by the considered CD setup, both have been manually geographically and geometrically aligned. The Sentinel-2 image used in this section is a $540 \times 525 \times 3$ image with 10m spatial resolution and composed of 3 spectral bands corresponding to visible RGB (Bands 2 to 4). On the other hand, Sentinel-1 reference image is a 540×525 interferometric wide swath high resolution ground range detected multi-looked SAR intensity image with a spatial resolution of 10m according to 5 looks in the range direction.

Generating the changes — Using a procedure similar to the one proposed by Ferraris et al. (2017a), given the reference images X_α^{ref} ($\alpha \in \{1, 2\}$), and a previously generated change mask $\mathbf{m} \in \mathbb{R}^{N_\alpha}$, a change image X_α^{ch} can be generated as

$$X_\alpha^{\text{ch}} = \vartheta(X_\alpha^{\text{ref}}, \mathbf{m}) \quad (36)$$

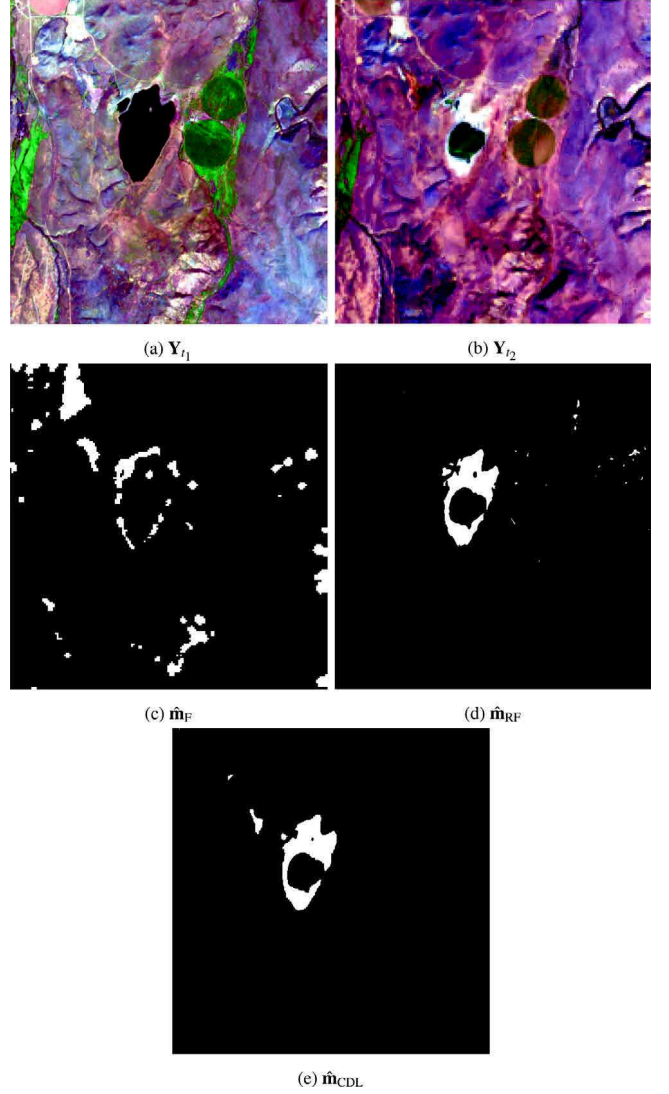


Fig. 6. Real images affected by real changes without ground truth, Scenario 1 (different spatial resolutions): (a) Sentinel-2 MS image Y_{t_1} acquired on 04/12/2016, (b) Landsat 8 MS image Y_{t_2} acquired on 09/22/2015, (c) change map \hat{m}_F of the fuzzy method, (d) change map \hat{m}_{RF} of the robust fusion method and (e) change map \hat{m}_{CDL} of the proposed method.

where the change-inducing functions $\vartheta : \mathbb{R}^{L \times N_\alpha} \times \mathbb{R}^{N_\alpha} \rightarrow \mathbb{R}^{L \times N_\alpha}$ is defined to simulate realistic changes in some pixels of the reference images. A set of 10 predefined change masks has been designed according to specific copy-paste change rules similar as the ones introduced by Ferraris et al. (2017a).

Generating the observed images — The observed images are generated under the previously defined 3 distinct scenarios involving 3 pairs of images, namely,

- Scenario 1 considers two optical images,
- Scenario 2 considers two SAR images,
- Scenario 3 considers a SAR image and an optical image.

Each test set pair $\{X_{\alpha_1}^{\text{ref}}, X_{\alpha_2}^{\text{ch}}\}$ is formed by considering $(\alpha_1, \alpha_2) = (\alpha, \alpha)$ with $\alpha = 1$ for Scenario 1 and $\alpha = 2$ for Scenario 2. Conversely, for Scenario 3 handling multimodal images, two test pairs can be formed considering $\alpha_1 \neq \alpha_2$, i.e., $(\alpha_1, \alpha_2) \in \{(1, 2), (2, 1)\}$.

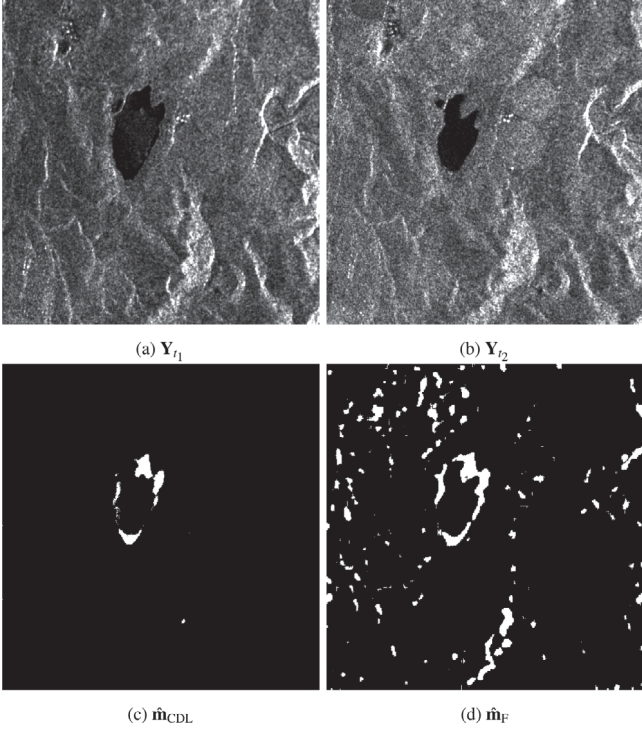


Fig. 7. Real images affected by real changes without ground truth, Scenario 2: (a) Sentinel-1 SAR image Y_{t_1} acquired on 04/12/2016, (b) Sentinel-1 SAR image Y_{t_2} acquired on 10/28/2016, (c) change map \hat{m}_F of the fuzzy method and (d) change map \hat{m}_{CDL} of the proposed method.

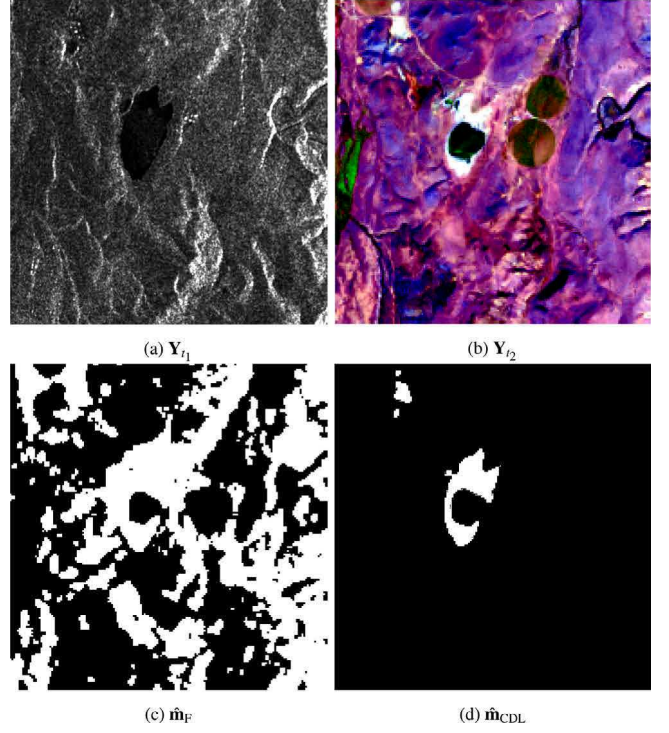


Fig. 9. Real images affected by real changes without ground truth, Scenario 3 (different spatial resolutions): (a) Sentinel-1 SAR image Y_{t_1} acquired on 04/12/2016, (b) Landsat 8 MS image Y_{t_2} acquired on 09/22/2015, (c) change map \hat{m}_F of the fuzzy method and (d) change map \hat{m}_{CDL} of the proposed method.

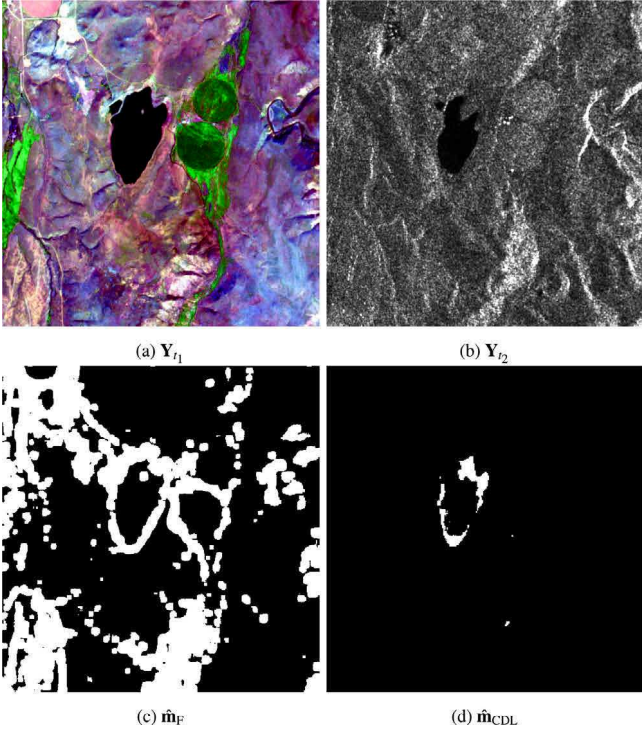


Fig. 8. Real images affected by real changes without ground truth, Scenario 3 (same spatial resolution): (a) Sentinel-2 MS image Y_{t_1} acquired on 04/12/2016, (b) Sentinel-1 SAR image Y_{t_2} acquired on 10/28/2016, (c) change map \hat{m}_F of the fuzzy method and (d) change map \hat{m}_{CDL} of the proposed method.

5.4.2. Results

The ROC curves displayed in Fig. 10 with corresponding metrics in Table 2 correspond to the CD results obtained for each specific scenario. These results are discussed below.

Scenario 1: optical vs. optical — The ROC curves displayed in Fig. 10(a) with corresponding metrics in Table 2 (first two rows) correspond to the CD results obtained from a pair of optical observed images. These results show that the robust fusion achieves the best CD performance. This method has the benefit of exploring the joint model between the optical images, contrary to the fuzzy and proposed methods. Nevertheless, the proposed method achieves very similar performance. More importantly, they provide almost perfect detection even for very low PFA, i.e., for very low energy changes. On the other hand, the fuzzy method suffers from non detection and false alarm, even when applying the iterative strategy with a similar parameter selection approach as in Gong et al. (2016). This happens mostly in low energy change regions. One possible explanation is that the iterative selection is not able to distinguish between low energy and unchanged pixels, which may bias the coupling of dictionaries. Also, the disjoint reconstruction cannot properly deal with low energy changes because coupling is not perfect. In addition, as the methods directly work with the observed images without estimating the latent image, noise could be interpreted as a change, thus increasing the false alarm rate.

Scenario 2: SAR vs. SAR — As in the previous case, this dual scenario considers homologous observed SAR images. In this case the ROC curves are displayed in Fig. 10(b) with corresponding metrics in Table 2 (3rd and 4th rows). Fig. 10(b) shows that the proposed method offers the highest precision among the compared methods and keeps a high level of detection compared to the Scenario 1. The fuzzy method presents a better accuracy result compared to optical images. One of the reasons is that optical images are generally characterized by richer information, which makes the dictionary coupling more difficult than

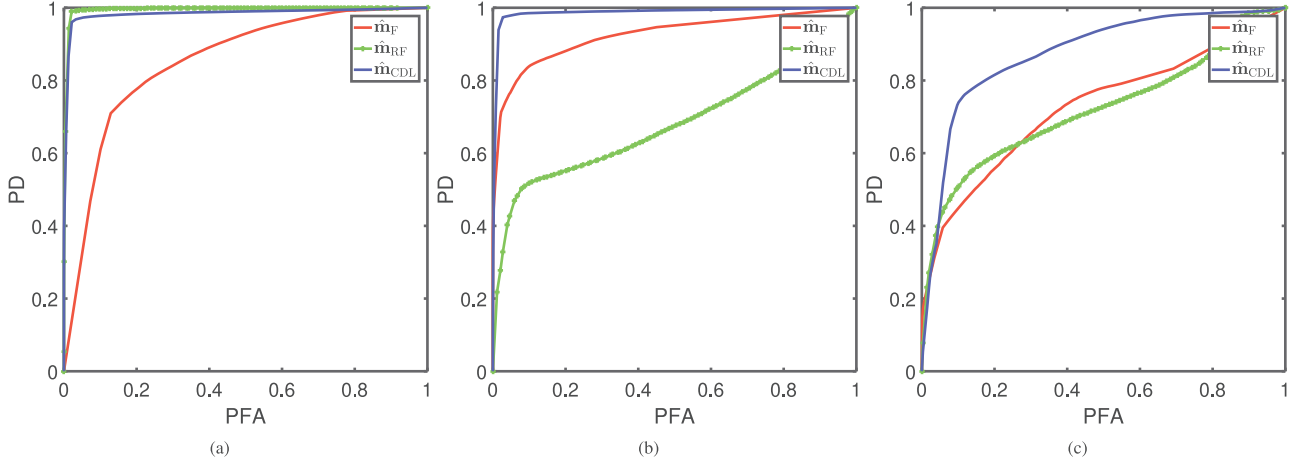


Fig. 10. Real images affected by synthetic changes: ROC curves for (a) Scenario 1, (b) Scenario 2, (c) Scenario 3.

Table 2

Real images affected by synthetic changes for Scenarios 1–3: quantitative detection performance (AUC and distance).

		\hat{m}_F	\hat{m}_{RF}	\hat{m}_{CDL}
Scenario 1	AUC	0.8520	0.9946	0.9838
	Dist.	0.7867	0.9802	0.9677
Scenario 2	AUC	0.9251	0.6819	0.9871
	Dist.	0.8587	0.6185	0.9727
Scenario 3	AUC	0.7277	0.7227	0.8755
	Dist.	0.6758	0.6604	0.8097

for two SAR images. At the end, the robust fusion CD method shows a very low detection accuracy as it is not suited to deal with SAR images.

Scenario 3: optical vs. SAR — This scenario corresponds to a more difficult problem than the previous one. The physical information extracted in each image cannot be directly related in the observational space, contrary to the previous scenarios. The ROC curve are displayed in Fig. 10(c) with corresponding metrics in Table 2 (last two rows). As in Scenario 2, Fig. 10(c) shows that the proposed method still offers the highest detection accuracy, while the other methods present a very poor performance. Regarding the fuzzy method, the dictionary and the subsequent sparse code estimations are severely affected by the differences in terms of dimensionality of measurements and dynamics. Even by tuning the algorithmic parameters to increase the weight of the image of lowest dynamics (or lowest resolution), the dictionaries are not properly coupled. Note that, to use the robust fusion method in this challenging scenario, a spectral degradation has been artificially applied to reach the same *spectral* resolution for the two images. This has been achieved by considering a band-averaging to finally form a panchromatic image. Resulting detection performance is even poorer than the fuzzy method because it supposes the same physical information between images. Only strong related changes are detected in this case.

5.5. Implementation details

This paragraph briefly discusses some computational aspects of the proposed method. Concerning the algorithmic implementation, the code was implemented in Matlab and run on a Windows platform equipped with a Intel Core i7 8GB RAM CPU. Depending on the size of the images and on the number of iterations, analyzing a pair of images as those considered in the experiments described in this section may require one hour. The computational bottleneck is the memory required to store parameters for each image and possibly large temporary byproduct variables, in particular those relying on large

matrix computation such as $A_\alpha A_\alpha^T$ or $D_\alpha^T D_\alpha$. The size of such matrices depends directly on the numbers N_1 and N_2 of pixels of the observed images and the chosen size N_d of the dictionaries. Note however that some computations could have been conducted in parallel to speed up the procedure, e.g., optimizing independently w.r.t. the dictionaries D_1 and D_2 (see Section 4.4), optimizing independently w.r.t. the latent images X_1 and X_2 (see Section 4.6), and optimizing w.r.t. the scaling matrix S independently from the coding matrix A_2 (see Section 4.5). Moreover, we experimentally noticed that the solution reached after very few iterations is generally highly satisfactory. Nevertheless, after more iterations, the shapes of the dictionary atoms visually seem to be much more representative and the sparseness of the code is significantly enforced.

Another important implementation aspect is the initialization of the optimization procedure, which is a critical issue because of the highly nonconvex nature of the problem. One may think of initializing the dictionary estimates by applying some techniques such as a (coupled) k-means clustering or the method proposed by Gong et al. (2016). Nevertheless, as also noticed by Seichepine et al. (2014), it was empirically observed that a better strategy is to randomly select coupled patches from the input data to form the initial dictionaries. This strategy may prevent the gradient to be initially stuck into a local minimum induced by weakly coupled dictionaries. As for the codes, since they are composed of a considerable number of block variables, one may pay attention to possibly exploding gradients due to high values. Initialization with zeros may induce the gradient to be stuck into local minima. Thus, one relevant strategy consists in initializing the codes with small random values.

6. Conclusion

This paper proposed an unsupervised multimodal change detection technique to handle the most common remote sensing imagery modalities. The technique was based on the definition of a pair of latent images related to the observed images through a direct observation model. These latent images were modeled thanks to a coupled dictionary and sparse codes which provide a common representation of the homologous patches in the latent image pair. The differences between estimated codes were assumed to be spatially sparse, implicitly locating the changes. Inferring these representations, as well as the latent images, was formulated as an inverse problem which was solved by the proximal alternate minimization iterative algorithm dedicated to nonsmooth and nonconvex functions. Contrary to the methods already proposed in the literature, scaling problems due to differences in resolutions and/or dynamics were solved by introducing a scaling matrix relating coupled atoms. A simulation protocol allowed the performance of the proposed technique in terms of detection and precision to be

assessed and compared with the performance of three algorithms. A real dataset collecting images from different multispectral and SAR sensors at the same region was used to assess the reliability of the proposed method. Results showed that the method outperformed all state-of-the-art comparable methods in multimodal scenarios while presenting similar results as methods benefiting from prior knowledge of the scenario modeling. Future works include considering more complex image statistical models, such as non-Gaussian distribution for optical images (Zanetti et al., 2015).

Acknowledgment

The authors would like to thank Prof. Jose M. Bioucas-Dias, Universidade de Lisboa, Portugal, for fruitful discussion regarding this work.

Appendix A. Data-fitting terms and corresponding proximal operators

The data-fitting term $D(\cdot|\cdot)$ is intimately related to the modality of the target image. This term defines the negative log-likelihood function relating the observed and latent images. Below, the most common data fitting terms and their associated proximal mappings are derived, defined as

$$\text{prox}_{D(\mathbf{Y}|\cdot)}^\eta(\mathbf{U}) = \underset{\mathbf{X}}{\text{argmin}} D(\mathbf{Y}|\mathbf{X}) + \frac{\eta}{2} \|\mathbf{X} - \mathbf{U}\|_F^2. \quad (\text{A.1})$$

A.1. Multiband optical images

Multiband optical images represent the most common modality of remotely sensed images. For this modality, the noise model may take into account several different noise sources (Deger et al., 2015). Nevertheless, it is commonly considered as additive Gaussian, up to some considerations in the acquisition, for instance sufficient number of the arriving photons. Therefore, the direct model $T_{\text{MO}}[\cdot]$ in (1) can be expressed as

$$\mathbf{Y} = \mathbf{X} + \mathbf{N} \quad (\text{A.2})$$

where the noise matrix \mathbf{N} is assumed to be distributed according to a matrix normal distribution (see, e.g., (Ferraris et al., 2017b) for more details). Consequently, by assuming the noise components are independent and identically distributed² (i.i.d.), the data-fitting term associated with multiband optical images is

$$D_{\text{MO}}(\mathbf{Y}|\mathbf{X}) = \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2. \quad (\text{A.3})$$

An explicit proximal operator associated with this function can be derived as

$$\text{prox}_{D_{\text{MO}}(\mathbf{Y}|\cdot)}^\eta(\mathbf{U}) = \frac{\mathbf{Y} + \eta\mathbf{U}}{\eta + 1} \quad (\text{A.4})$$

A.2. Multi-look intensity synthetic aperture radar images

SAR images correspond to the second most common modality of remote sensing images used in many applications. One of the main characteristics of such modality is that it allows to measure the scene in poor weather conditions and also during the night since SAR is an active sensor. Nevertheless, this configuration yields the speckle phenomenon, resulting from random fluctuations of the reflectivity of the backscattered signals. Many studies have been conducted to understand and mitigate the speckle phenomenon. A common approach that helps to decrease the speckle level while increasing the SNR consists in

averaging samples of the same pixel acquired over independent observations. This procedure is usually referred to as multi-look processing. According to this strategy, the generative model is considered as a multiplicative perturbation by i.i.d random variables $\mathbf{N} = [n_1, \dots, n_N]$ following a common gamma probability density function in intensity images with unit mean $E[n_i] = 1$ and variance $\text{var}[n_i] = \frac{1}{r}$ where r is the number of looks. The direct model $T_{\text{SAR}}[\cdot]$ can thus be written as

$$\mathbf{Y} = \mathbf{X} \odot \mathbf{N} \quad (\text{A.5})$$

where \odot denotes the termwise (i.e., Hadamard) product.

By assuming pixel independence, the data-fitting term for each pixel can be expressed as the sum of Itakura–Saito divergences

$$D_{\text{SAR}}(\mathbf{Y}|\mathbf{X}) = \sum_{i=1}^N \left(\frac{y_i}{x_i} - \log \frac{y_i}{x_i} - 1 \right) \quad (\text{A.6})$$

This function has been widely considered for speckle removing (Aubert and Aujol, 2008; Woo and Yun, 2013) and also music analysis (Févotte et al., 2009). Nevertheless, it usually leads to a challenging non-convex problem which admits more than one global solution. In Sun and Févotte (2014), the associated proximal operator is derived by computing the root of a 3rd degree-polynomial equation. An alternative consists in considering an approximation by resorting to a log-transform of the data, e.g., leading to an I-divergence (Woo and Yun, 2013; Steidl and Teuber, 2010). Up to a constant, this divergence can be rewritten equivalently as a Kullback–Leibler divergence which is closely related to Poisson modeling (Figueiredo and Bioucas-Dias, 2010)

$$D_{\text{SAR}}(\mathbf{Y}|\mathbf{X}) = \sum_{i=1}^N (x_i - y_i \log x_i). \quad (\text{A.7})$$

This data-fitting term leads to an explicit proximal operator for the i th component given by

$$\text{prox}_{D_{\text{SAR}}(\cdot|\cdot)}^\eta(u_i) = \frac{1}{2} \left(u_i - \frac{1}{\eta} + \sqrt{\left(u_i - \frac{1}{\eta} \right)^2 + \frac{4y_i}{\eta}} \right). \quad (\text{A.8})$$

Appendix B. Usual proximal mappings involved in the parameter updates

The projections and proximal operators involved on PALM algorithm (Bolte et al., 2014) and described in Algorithm 1 are properly defined as:

- The proximal map for \mathbf{A}_1 accounting for the sum $\lambda \|\cdot\|_1 + \iota_{\geq 0}(\cdot)$ is explicitly given by:

$$\text{prox}_{\lambda \|\cdot\|_1 + \geq 0}^\eta(\mathbf{A}_1) = \max \left(|a_{1,(j,i)}| - \frac{\lambda}{\eta}, 0 \right) \quad \forall (i, j) \quad (\text{B.1})$$

- The proximal map for $\Delta \mathbf{A}$ accounting for the $\gamma \|\cdot\|_{2,1}$ is explicitly given by:

$$\text{prox}_{\gamma \|\cdot\|_{2,1}}^\eta(\Delta \mathbf{A}) = \begin{cases} \left(1 - \frac{\gamma}{\eta \|\Delta \mathbf{a}_i\|_2} \right) \Delta \mathbf{a}_i & \text{if } \|\Delta \mathbf{a}_i\|_2 > \frac{\gamma}{\eta} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B.2})$$

- Projecting \mathbf{D} onto set S can be computed explicitly based on Thouvenin (2017), Bolte et al. (2014) which is given by:

$$\mathcal{P}_S(\mathbf{D}) = \frac{\mathcal{P}_+(\mathbf{d}_i)}{\|\mathcal{P}_+(\mathbf{d}_i)\|_2} \quad \forall i = 1 \dots N_d \quad (\text{B.3})$$

with

$$\mathcal{P}_+(\mathbf{d}_i) = \max(0, d_{(j,i)}) \quad \forall j = 1 \dots L \quad (\text{B.4})$$

- Projecting \mathbf{S} onto set C is explicitly given by:

$$\mathcal{P}_C(\mathbf{S}) = \begin{cases} \max(0, s_{(j,i)}) & \forall i = j \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.5})$$

² Pixelwise independence of the noise is a common assumption while spectral whiteness of the noise can be ensured by applying a whitening transform as pre-processing.

References

- Aharon, M., Elad, M., Bruckstein, A., 2006. K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* 54 (11), 4311–4322.
- Alberga, V., Idrissa, M., Lacroix, V., Inglada, J., 2007. Comparison of similarity measures of multi-sensor images for change detection applications. In: *IEEE Int. Conf. Geosci. Remote Sens., IGARSS. IEEE*, pp. 2358–2361.
- Aubert, G., Aujol, J.-F., 2008. A variational approach to removing multiplicative noise. *SIAM J. Appl. Math.* 68 (4), 925–946.
- Bach, F., 2011. Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.* 4 (1), 1–106.
- Bolte, J., Sabach, S., Teboulle, M., 2014. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Math. Program.* 146 (1–2), 459–494.
- Bovolo, F., Bruzzone, L., 2007. A theoretical framework for unsupervised change detection based on change vector analysis in the polar domain. *IEEE Trans. Geosci. Remote Sens.* 45 (1), 218–236.
- Bovolo, F., Bruzzone, L., 2015. The time variable in data fusion: a change detection perspective. *IEEE Geosci. Remote Sens. Mag.* 3 (3), 8–26.
- Brunner, D., Lemoine, G., Bruzzone, L., 2010. Earthquake damage assessment of buildings using VHR optical and SAR imagery. *IEEE Trans. Geosci. Remote Sens.* 48 (5), 2403–2420.
- Bruzzone, L., Prieto, D.F., Serpico, S.B., 1999. A neural-statistical approach to multi-temporal and multisource remote-sensing image classification. *IEEE Trans. Geosci. Remote Sens.* 37 (3), 1350–1359.
- Campbell, J.B., Wynne, R.H., 2011. *Introduction to Remote Sensing*, fifth ed. Guilford Press, New York.
- Cavalcanti, Y.C., Oberlin, T., Dobigeon, N., Tauber, C., 2017. Unmixing dynamic PET images with a PALM algorithm. In: *Proc. European Signal Process. Conf., EUSIPCO. IEEE*, pp. 425–429.
- Chabert, M., Tourneret, J.-Y., Poulain, V., Inglada, J., 2010. Logistic regression for detecting changes between databases and remote sensing images. In: *Proc. IEEE Int. Conf. Geosci. Remote Sens., IGARSS. IEEE*, pp. 3198–3201.
- Chambolle, A., 2004. An algorithm for total variation minimization and applications. *J. Math. Imaging Vis.* 20 (1–2), 89–97.
- Chen, S.S., Donoho, D.L., Saunders, M.A., 2001. Atomic decomposition by basis pursuit. *SIAM Rev.* 43 (1), 129–159.
- Combettes, P.L., Wajs, V.R., 2005. Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.* 4 (4), 1168–1200.
- Coppin, P., Jonckheere, I., Nackaerts, K., Muys, B., Lambin, E., 2004. Review article digital change detection methods in ecosystem monitoring: a review. *Int. J. Remote Sens.* 25 (9), 1565–1596.
- Deger, F., Mansouri, A., Pedersen, M., Hardeberg, J.Y., Voisin, Y., 2015. A sensor data based denoising framework for hyperspectral images. *Opt. Express* 23 (3), 1938.
- Elad, M., Aharon, M., 2006. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. Image Process.* 15 (12), 3736–3745.
- Eriksson, K., Estep, D., Johnson, C., 2004. Lipschitz continuity. In: *Applied Mathematics: Body and Soul*, vol. 1. Springer-Verlag Berlin Heidelberg.
- European Space Agency, 2017a. Sentinel-1. <http://www.esa.int>.
- European Space Agency, 2017b. Sentinel-2. <http://www.esa.int>.
- Feng, W., Sui, H., Tu, J., Huang, W., Xu, C., Sun, K., 2018. A novel change detection approach for multi-temporal high-resolution remote sensing images based on rotation forest and coarse-to-fine uncertainty analyses. *Remote Sens.* 10 (7), 1077.
- Ferraris, V., Dobigeon, N., Wei, Q., Chabert, M., 2017a. Detecting changes between optical images of different spatial and spectral resolutions: a fusion-based approach. *IEEE Trans. Geosci. Remote Sens.* 1–13.
- Ferraris, V., Dobigeon, N., Wei, Q., Chabert, M., 2017b. Robust fusion of multiband images with different spatial and spectral resolutions for change detection. *IEEE Trans. Comput. Imag.* 3 (2), 175–186.
- Févotte, C., Bertin, N., Durrieu, J.-L., 2009. Nonnegative matrix factorization with the Itakura-Saito Divergence: with application to music analysis. *Neural Comput.* 21 (3), 793–830.
- Févotte, C., Dobigeon, N., 2015. Nonlinear hyperspectral unmixing with robust nonnegative matrix factorization. *IEEE Trans. Image Process.* 24 (12), 4810–4819.
- Figueiredo, M., Bioucas-Dias, J., 2010. Restoration of poissonian images using alternating direction optimization. *IEEE Trans. Image Process.* 19 (12), 3133–3145.
- Fountoulakis, K., Gondzio, J., 2016. A second order method for strongly convex ℓ_1 -regularization problems. *Math. Program.* 156 (1–2), 189–219.
- Gelman, A. (Ed.), 2004. *Bayesian Data Analysis*, second ed. *Texts in Statistical Science*, Chapman & Hall/CRC, Boca Raton, Fla.
- Gong, M., Zhang, P., Su, L., Liu, J., 2016. Coupled dictionary learning for change detection from multisource data. *IEEE Trans. Geosci. Remote Sens.* 54 (12), 7077–7091.
- Huang, X., Yang, W., Xia, G., Liao, M., 2015. Superpixel-based change detection in high resolution sar images using region covariance features. In: *2015 8th International Workshop on the Analysis of Multitemporal Remote Sensing Images, (Multi-Temp)*, pp. 1–4.
- Inglada, J., 2002. Similarity measures for multisensor remote sensing images. In: *Proc. IEEE Int. Conf. Geosci. Remote Sens., IGARSS. IEEE*, pp. 104–106.
- Jensen, T.L., Jørgensen, J.H., Hansen, P.C., Jensen, S.H., 2012. Implementation of an optimal first-order method for strongly convex total variation regularization. *BIT Numer. Math.* 52 (2), 329–356.
- Kawamura, J.G., 1971. Automatic recognition of changes in urban development from aerial photographs. *IEEE Trans. Syst., Man, Cybern.* SMC-1 (3), 230–239.
- Lu, D., Mausel, P., Brondizio, E., Moran, E., 2004. Change detection techniques. *Int. J. Remote Sens.* 25 (12), 2365–2401.
- Lu, X., Yuan, Y., Zheng, X., 2017. Joint dictionary learning for multispectral change detection. *IEEE Trans. Cybern.* 47 (4), 884–897.
- Ma, L., Moisan, L., Yu, J., Zeng, T., 2013. A dictionary learning approach for Poisson image deblurring. *IEEE Trans. Med. Imag.* 32 (7), 1277–1289.
- Mairal, J., 2014. Sparse modeling for image and vision processing. *Found. Trends Comput. Graph. Vis.* 8 (2–3), 85–283.
- Mairal, J., Bach, F., Ponce, J., Sapiro, G., 2009. Online dictionary learning for sparse coding. In: *Proc. Int. Conf. Machine Learning, ICML. ACM*, pp. 689–696.
- Mallat, S., 2009. *A Wavelet Tour of Signal Processing: The Sparse Way*, third ed. Academic Press, Boston.
- Mercier, G., Moser, G., Serpico, S., 2008. Conditional copulas for change detection in heterogeneous remote sensing images. *IEEE Trans. Geosci. Remote Sens.* 46 (5), 1428–1441.
- Olshausen, B.A., Field, D.J., 1997. Sparse coding with an overcomplete basis set: a strategy employed by V1. *Vis. Res.* 37, 3311–3325.
- Parikh, N., Boyd, S., et al., 2014. Proximal algorithms. *Found. Trends Optim.* 1 (3), 127–239.
- Prendes, J., Chabert, M., Pascal, F., Giros, A., Tourneret, J.-Y., 2015. A new multivariate statistical model for change detection in images acquired by homogeneous and heterogeneous sensors. *IEEE Trans. Image Process.* 24 (3), 799–812.
- Seichepine, N., Essid, S., Févotte, C., Cappe, O., 2014. Soft nonnegative matrix co-factorization. *IEEE Trans. Signal Process.* 62 (22), 5940–5949.
- Singh, A., 1989. Review article digital change detection techniques using remotely-sensed data. *Int. J. Remote Sens.* 10 (6), 989–1003.
- Solano-Correa, Y.T., Bovolo, F., Bruzzone, L., 2018. An approach for unsupervised change detection in multitemporal vhr images acquired by different multispectral sensors. *Remote Sens.* 10 (4), 601–620.
- Solberg, A.H.S., Taxt, T., Jain, A.K., 1996. A Markov random field model for classification of multisource satellite imagery. *IEEE Trans. Geosci. Remote Sens.* 34 (1), 100–113.
- Steidl, G., Teuber, T., 2010. Removing multiplicative noise by Douglas-Rachford splitting methods. *J. Math. Imaging Vision* 36 (2), 168–184.
- Sun, D.L., Févotte, C., 2014. Alternating direction method of multipliers for non-negative matrix factorization with the beta-divergence. *J. Math. Imaging Vis.* 6201–6205.
- Thouvenin, P.-A., 2017. *Modeling spatial and temporal variabilities in hyperspectral image unmixing* (Ph.D. thesis). Institut National Polytechnique de Toulouse, Toulouse.
- Thouvenin, P.-A., Dobigeon, N., Tourneret, J.-Y., 2016. Online unmixing of multitemporal hyperspectral images accounting for spectral variability. *IEEE Trans. Image Process.* 25 (9), 3979–3990.
- United States Geological Survey, 2017. Landsat 8. <https://landsat.usgs.gov/landsat-8>.
- Woo, H., Yun, S., 2013. Proximal linearized alternating direction method for multiplicative denoising. *SIAM J. Sci. Comput.* 35 (2), B336–B358.
- Wright, S., Nowak, R., Figueiredo, M., 2009. Sparse reconstruction by separable approximation. *IEEE Trans. Signal Process.* 57 (7), 2479–2493.
- Yang, J., Wang, Z., Lin, Z., Cohen, S., Huang, T., 2012. Coupled dictionary training for image super-resolution. *IEEE Trans. Image Process.* 21 (8), 3467–3478.
- Yang, J., Wright, J., Huang, T.S., Ma, Y., 2010. Image super-resolution via sparse representation. *IEEE Trans. Image Process.* 19 (11), 2861–2873.
- Zanetti, M., Bovolo, F., Bruzzone, L., 2015. Rayleigh-rice mixture parameter estimation via EM algorithm for change detection in multispectral images. *IEEE Trans. Image Process.* 24 (12), 5004–5016.
- Zeyde, R., Elad, M., Protter, M., 2010. On single image scale-up using sparse-representations. In: *International Conference on Curves and Surfaces*. Springer, pp. 711–730.