



CD-GAN: A robust fusion-based generative adversarial network for unsupervised remote sensing change detection with heterogeneous sensors

Jin-Ju Wang, Nicolas Dobigeon, Marie Chabert, Ding-Cheng Wang, Ting-Zhu Huang, Jie Huang

► To cite this version:

Jin-Ju Wang, Nicolas Dobigeon, Marie Chabert, Ding-Cheng Wang, Ting-Zhu Huang, et al.. CD-GAN: A robust fusion-based generative adversarial network for unsupervised remote sensing change detection with heterogeneous sensors. *Information Fusion*, 2024, 107, pp.102313. 10.1016/j.inffus.2024.102313 . hal-04483118

HAL Id: hal-04483118

<https://hal.science/hal-04483118>

Submitted on 29 Feb 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

CD-GAN: a robust fusion-based generative adversarial network for unsupervised remote sensing change detection with heterogeneous sensors

Jin-Ju Wang^{a,*}, Nicolas Dobigeon^b, Marie Chabert^b, Ding-Cheng Wang^c, Ting-Zhu Huang^c, Jie Huang^c

^a*School of Mathematics, Southwest Minzu University, Chengdu, Sichuan 610041, China*

^b*University of Toulouse, IRIT/INP-ENSEEIH Toulouse, BP 7122, 31071 Toulouse Cedex 7, France*

^c*School of Mathematical Sciences, University of Electronic Science and Technology of China (UESTC), Chengdu, Sichuan, China*

Abstract

In the context of Earth observation, change detection boils down to comparing images acquired at different times by sensors of possibly different spatial and/or spectral resolutions or different modalities (e.g., optical or radar). Even when considering only optical images, this task has proven to be challenging as soon as the sensors differ by their spatial and/or spectral resolutions. This paper proposes a novel unsupervised change detection method dedicated to images acquired by such so-called heterogeneous optical sensors. It capitalizes on recent advances which formulate the change detection task into a robust fusion framework. Adopting this formulation, the work reported in this paper shows that any off-the-shelf network trained beforehand to fuse optical images of different spatial and/or spectral resolutions can be easily complemented with a network of the same architecture and embedded into an adversarial framework to perform change detection. A comparison with state-of-the-art change detection methods demonstrates the versatility and the effectiveness of the proposed approach.

Keywords: heterogeneous sensors, change detection, adversarial training, fusion, spatial sparsity.

1. Introduction

In the context of Earth observation, change detection (CD) aims at identifying spatial areas that have been altered between two remote sensing image acquisitions [45, 4]. With the rapid development of urban areas and the frequent occurrence of natural disasters like floods and earthquakes, CD has become an essential task to monitor land cover evolution [14, 16, 48, 31]. To be effective, the CD methods have to account for the intrinsic characteristics of the two sensors acquiring the pair of images and their possible dissimilarities. When the sensors share the same modality and have the same spatial and spectral resolutions, the images they provide are usually referred to as homogeneous. Conversely, when the sensors differ by their modality or resolutions, the images are referred to as heterogeneous.

Due to the rapid development of remote sensing technologies, there is a growing availability of images coming from heterogeneous sensors. This offers new opportunities to perform

timely CD, which may be valuable in emergency situations. In this context, this paper specifically focuses on the challenging CD applicative scenario that involves optical sensors with different spatial and/or spectral resolutions. An archetypal scenario, considered in this paper for illustrative purpose and generally referred to as *complementary acquisition*, involves two optical images of complementary resolutions, i.e., a high spatial and low spectral resolution (HRLS) image and a low spatial and high spectral resolution (LRHS) image. Adopting a conventional dichotomy over the optical sensors, this could be *i)* a pair composed of a panchromatic (PAN) image and a multispectral (MS) image, *ii)* a pair composed of a PAN image and a hyperspectral (HS) image or *iii)* a pair composed of a MS image and an HS image.

The pipeline of most CD methods consists of three steps: image pre-processing, change image (CI) generation and binary change map (CM) generation [5]. In particular, when dealing with heterogeneous sensors, the pre-processing step aims at allowing a reliable pixel-by-pixel comparison to be performed between the two acquired images. The second step then generates the CI by comparing the pre-processed images pixel-by-pixel. The last step performs high level processing such as segmentation or classification on the CI to label the altered pixels and/or group of pixels.

When facing dissimilar spatial and/or spectral resolutions, the first step of the CD pipeline outlined above mainly rely on individual processing applied to each acquired image. It produces a pair of virtually observed images with the same spatial and spectral resolutions where two homologous pixels (e.g., with the same coordinates) depict the same spatial location in

*Part of this work has been supported by the ANR-3IA Artificial, Natural Intelligence Toulouse Institute (ANITI) under grant agreement ANITI ANR-19-P3IA-0004, NSFC (Grant No. 12171072, Grant No. 12371517), Key Projects of Applied Basic Research in Sichuan Province (Grant No. 2020YJ0216), National Key Research and Development Program of China (Grant No. 2020YFA0714001), and Fundamental Research Funds for the Central Universities (Grant No. ZYN2024018).

*Corresponding author

Email addresses: jinjuwang123@163.com (Jin-Ju Wang), nicolas.dobigeon@enseeiht.fr (Nicolas Dobigeon), marie.chabert@enseeiht.fr (Marie Chabert), wangdc@uestc.edu.cn (Ding-Cheng Wang), tingzhuhuang@126.com (Ting-Zhu Huang), huangjie_uestc@uestc.edu.cn (Jie Huang)

the observed scene. This step can be achieved following various strategies, such as filtering, interpolation, resampling, deconvolution or super-resolution [24, 13, 36, 59, 22, 2, 1]. For instance, a virtual HRLS image can be obtained by spatially interpolating and spectrally filtering an observed LRHS image. Another possibility is to spatially downsample the observed HRLS image and spectrally filtering the LRHS image to produce a pair of virtual images with the same low spatial and low spectral resolutions (LRLS). Thereby, the two acquired images are embedded into the same feature space where their virtual counterparts can be compared pixel-by-pixel [35, 34]. Instead of applying these transformations in the image domain, one alternative may consist in deriving a patch-based graph representation of each image. The structures of the two resulting graphs are then compared to locate changes [42, 43, 41].

However, the above pre-processing strategies are implemented on each acquired image individually. Thus they do not fully exploit their interdependence and complementarity in terms of spatial and spectral information. Moreover, some of these pre-processing consist in artificially degrading the spatial and/or spectral resolutions of at least one of the acquired images, obviously resulting in a loss of information. Instead of individually pre-processing the two acquired images, a general CD framework, initially introduced in [11] and later extended in [12, 10], has envisioned the CD problem from the prism of multiband image fusion. Interestingly, this task shares some key characteristics with the problem of CD between heterogeneous but complementary images. It aims at recovering a latent (i.e., unobserved or virtual) image by merging the spatial and spectral information brought by a pair of observed images with complementary resolutions [27, 50, 7, 54, 56, 49, 17, 23]. Tackling the CD task under the fusion paradigm, the spatial and spectral information of the two observed images is fully preserved, which enables the subsequent recovery of a CD map with the highest spatial and spectral resolutions. This strategy has demonstrated its superiority over state-of-the-art CD methods that mostly rely on feature space embedding.

Meanwhile, some recent works exploit the versatility of deep neural networks to perform CD in case of homogeneous or heterogeneous sensors [40, 26, 59, 25, 55]. For instance, the so-called slow feature analysis proposed by Du et al. [8] consists in comparing the features learnt by two deep networks, each dedicated to one of the observed image. Again, this strategy does not fully exploit the information shared by the two images since the features are extracted from the two images individually. In [28], cycle-consistent adversarial networks (CycleGANs) learn a subimage-to-subimage mapping, that allows the LRHS image to be embedded into a HRLS image space. This image translation can be also achieved thanks to particular network architectures embedding transformers to finally generate the CI [53]. However, most of these deep learning-based CD methods face several major shortcomings. Firstly, very few consider the case of heterogeneous acquisitions and most of the methods would need to be carefully adapted to handle images of different spatial and/or spectral resolutions. Secondly, the deep architectures solving the image translation problem hardly preserve the whole spatial and spectral information brought by the pair of

images, which may affect the objects and backgrounds. Thirdly, many of these deep architectures are primarily designed for supervised CD, i.e., their training requires image pairs taken before and after changes, with annotated changes.

This paper proposes an unsupervised CD method for heterogeneous sensors (e.g., handling a HRLS image and a LRHS image) by capitalizing on the fusion-based CD framework mentioned above. This framework has the great advantage of jointly exploiting the spatial and spectral information brought by the two observed images, contrary to most of the approaches discussed above which apply individual pre-processing. Moreover it leverages the exceptional performance reached by recent neural fusion architectures, without requiring image pairs with annotated changes for training. More precisely, this paper shows that any off-the-shelf pre-trained deep network designed to fuse heterogeneous images can be reused as part of an adversarial architecture to perform unsupervised CD. This fusion network, whose design can be left to the final user, can be systematically complemented with a network of the same architecture to estimate a CI image. When compared to more conventional approaches, the strategy adopted in this paper appears to be more flexible while competing favorably. Moreover, the proposed framework can be kept up-to-date by embedding even more powerful pre-trained fusion networks, taking benefit from expected future advances in that domain.

The main contributions of this paper can be summarized as follows: *i)* after casting the heterogeneous CD problem into a robust fusion framework, we show that this framework can be interpreted as an generative adversarial model, whose generator is composed of two essential building blocks, *ii)* since one of this block boils down to a fusion procedure, we suggest that any off-the-shelf fusion algorithm, such as a pre-trained deep neural network, can be easily complemented by a deep neural network to build an adversarial architecture dedicated to CD and *iii)* exploiting the property of fusion consistency, a specific loss function is derived to allow the overall architecture to be trained within an unsupervised adversarial framework.

This article is organized as follows. Section 2 recalls how the problem of CD can be cast as a robust fusion task. Capitalizing on this formulation, an adversarial framework is proposed in Section 3 and a possible architecture is detailed for the particular case of complementary acquisitions. Experimental results obtained on simulated datasets are reported in Section 4 to assess the efficiency, the versatility and the robustness of the proposed approach when compared to alternative CD methods. Section 5 illustrates the approach relevance when analyzing real datasets. Finally, Section 6 concludes the paper.

2. Background: robust fusion-based change detection

This paper focuses on the problem of detecting changes between two optical images denoted $\mathbf{Y}_1 \in \mathbb{R}^{m_1 \times n_1}$ and $\mathbf{Y}_2 \in \mathbb{R}^{m_2 \times n_2}$ acquired over the same scene, at different times t_1 and t_2 by two sensors \mathbf{S}_1 and \mathbf{S}_2 , respectively. Here, m_i and n_i stand for the numbers of bands and pixels of the image \mathbf{Y}_i ($i \in \{1, 2\}$), respectively. In our case of interest, these two images are assumed

to have different spatial and/or spectral resolutions, i.e., $n_1 \neq n_2$ and/or $m_1 \neq m_2$.

Remark (Complementary acquisitions) *Throughout this paper, we will consider the particular case of images with complementary resolutions, as in [12]. This choice is made for illustrative purposes and does not sacrifice the generality of the approach. This case can be stated by the following relations*

$$m_1 \geq m_2 \text{ and } n_2 \geq n_1. \quad (1)$$

In other words, \mathbf{Y}_1 and \mathbf{Y}_2 are LRHS and HRLS images respectively.

Because of the different spatial and/or spectral resolutions between \mathbf{Y}_1 and \mathbf{Y}_2 , a naive pixel-wise comparison of the two images cannot be conducted to locate the changes. To overcome this difficulty, the adversarial strategy proposed in this paper builds upon the robust fusion framework proposed in [12] and generalized in [10]. This fusion framework relates the two observed images \mathbf{Y}_1 and \mathbf{Y}_2 and two latent images $\mathbf{X}_1 \in \mathbb{R}^{m \times n}$ and $\mathbf{X}_2 \in \mathbb{R}^{m \times n}$ which share the same high spatial and high spectral resolutions with $n \geq \max\{n_1, n_2\}$ and $m \geq \max\{m_1, m_2\}$. More precisely, the latent and acquired images are related through the observation models

$$\mathbf{Y}_1 = \mathcal{H}_1(\mathbf{X}_1) \quad (2a)$$

$$\mathbf{Y}_2 = \mathcal{H}_2(\mathbf{X}_2) \quad (2b)$$

where $\mathcal{H}_1 : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m_1 \times n_1}$ and $\mathcal{H}_2 : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m_2 \times n_2}$ stand for degradation operators. These operators are assumed to be linear, i.e., $\mathcal{H}_i(\mathbf{A} + \lambda\mathbf{B}) = \mathcal{H}_i(\mathbf{A}) + \lambda\mathcal{H}_i(\mathbf{B})$ for $i \in \{1, 2\}$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{B} \in \mathbb{R}^{m \times n}$ and $\lambda \in \mathbb{R}$.

Remark (Complementary acquisitions) *In case of complementary acquisitions, the operators $\mathcal{H}_1(\cdot)$ and $\mathcal{H}_2(\cdot)$ stand for spatial and spectral degradations respectively. As in [12] and in most of previous works from the literature dealing with multi-band imaging fusion [52, 7, 54, 49, 17], they are defined as*

$$\mathcal{H}_1(\mathbf{X}_1) \approx \mathbf{X}_1 \mathbf{R},$$

$$\mathcal{H}_2(\mathbf{X}_2) \approx \mathbf{L} \mathbf{X}_2.$$

Here, the matrix \mathbf{R} is generally decomposed into $\mathbf{R} = \mathbf{B}\mathbf{S}$, where \mathbf{B} stands for a spatially-invariant blurring operation and \mathbf{S} corresponds to a regular subsampling. The matrix \mathbf{L} is a spectral degradation specified by the sensor spectral response. The symbol \approx accounts for potential mis modeling or measurement noise. The latent space $\mathbb{R}^{m \times n}$ embeds images with spectral and spatial resolutions defined by the LRHS and HRLS images respectively, i.e., $m = m_1$ and $n = n_2$.

It is worth noting that the latent images \mathbf{X}_1 and \mathbf{X}_2 share the same spatial and spectral resolutions. As a consequence, these two latent images could be pixel-wisely compared to detect any changes that may have occurred between the two acquisition times. Based on this finding, CD can be reframed as an inference problem referred to as robust fusion. The resolution of this problem can be decomposed into two main steps, which are detailed below.

2.1. CI inference

Recall that the two latent images \mathbf{X}_1 and \mathbf{X}_2 share the same spatial and spectral resolutions. If they were available, a CI denoted $\Delta\mathbf{X} \in \mathbb{R}^{m \times n}$ of high spatial and high spectral resolutions could be easily derived by a pixel-by-pixel difference. More precisely, this change image is computed as

$$\Delta\mathbf{X} = \mathbf{X}_2 - \mathbf{X}_1 \quad (3)$$

with $\Delta\mathbf{X} = [\Delta\mathbf{x}_1, \dots, \Delta\mathbf{x}_n]$, where $\Delta\mathbf{x}_i$ is the $m \times 1$ change vector associated to the pixel $i \in \{1, \dots, n\}$.

Then, conventional CD methods dedicated to homogeneous sensors, such as change vector analysis (CVA) and its extensions [3, 9], can be used to derive the binary CM from the CI. In its canonical formulation, CVA consists in computing the pixel-wise energy $\mathbf{e} = [e_1, \dots, e_n] \in \mathbb{R}^n$ of the change image where, for $i = 1, \dots, n$,

$$e_i = \|\Delta\mathbf{x}_i\|_2.$$

The components e_i of \mathbf{e} with the smallest (resp. highest) values most likely correspond to unchanged (resp. changed) pixels, whose indices are gathered in the set Ω_u (resp. Ω_c). Thus, a natural decision rule consists in thresholding the energy change image, i.e., for $i = 1, \dots, n$

$$i \in \begin{cases} \Omega_c & \text{if } e_i > \tau, \\ \Omega_u & \text{otherwise.} \end{cases} \quad (4)$$

Here, the threshold τ adjusts the trade-off between the probability of false alarm and the probability of detection. It can be chosen using a dedicated method, as in [39]. Then, the final binary CM $\mathbf{d} = [d_1, \dots, d_n] \in \{0, 1\}^n$ can be derived as

$$d_i = \begin{cases} 1 & \text{if } i \in \Omega_c, \\ 0 & \text{if } i \in \Omega_u. \end{cases} \quad (5)$$

It is worth noting that, as stated in many works such as [43], the changes are generally related to any evolution of the land cover, e.g., due to natural disasters, seasonal trends or human induced modifications. Conversely, when dealing with optical images, variations due to illumination should not be identified as changes since they result from the acquisition conditions. Under this paradigm, the conventional formulation of CD assumes that most of the pixels are expected to be unaffected by changes. In other words, CD is a particular instance of an unbalanced binary classification problem since it consists in recovering the binary CM \mathbf{d} with $\text{card}\{\Omega_u\} \gg \text{card}\{\Omega_c\}$. Thus, the threshold τ should be adjusted to identify the most relevant changes in the CI, in particular those that cannot be explained only by illumination variations.

2.2. Fusion

Reciprocally, if the CI $\Delta\mathbf{X}$ was known, the latent images \mathbf{X}_1 and \mathbf{X}_2 could be inferred by solving a multiband image fusion problem. Indeed, adopting the framework introduced in [12]

and exploiting the identity (3) and the linearity of the degradation operators, the direct models (2) can be equivalently rewritten as

$$\mathbf{Y}_1 = \mathcal{H}_1(\mathbf{X}_1), \quad (6a)$$

$$\tilde{\mathbf{Y}}_2 = \mathcal{H}_2(\mathbf{X}_1) \quad (6b)$$

where¹

$$\tilde{\mathbf{Y}}_2 \triangleq \mathbf{Y}_2 - \mathcal{H}_2(\Delta\mathbf{X}), \quad (7)$$

is the so-called *corrected* image, i.e., the image that would be acquired by \mathbf{S}_2 at t_1 (before a change occurs). From the set of equations (6), it clearly appears that recovering the latent image \mathbf{X}_1 from the observed image \mathbf{Y}_1 and the corrected image $\tilde{\mathbf{Y}}_2$ is an image fusion task. The other latent image \mathbf{X}_2 can be subsequently derived from \mathbf{X}_1 and $\Delta\mathbf{X}$ using (3).

Remark (Complementary acquisitions) *Fusing a pair of images resulting from complementary acquisitions aims at recovering a latent image whose spatial and spectral resolutions are the highest of the observed image ones. This problem has motivated a wide bunch of research works over the three last decades, specifically dealing with pansharpening when fusing PAN and MS images [46], hyperspectral pansharpening when fusing PAN and HS images [29, 6] or multiband image fusion when fusing MS and HS images [51, 7, 23].*

2.3. From fusion consistency to CD

From what precedes, heterogeneous CD can be decomposed into two entangled steps, namely CI inference and fusion. These two steps cannot be performed independently since each step relies on key quantities provided as outputs by the other step. To overcome this difficulty, the robust fusion-based CD approach leverages on a consistency property fulfilled by any fusion method. This property states that a well-designed fusion process should be reversible. This means that the observed image \mathbf{Y}_1 is expected to be as close as possible to the image that would be obtained by applying the acquisition process to the estimated fused image $\hat{\mathbf{X}}_1$ [47]. In other words, the so-called *predicted* image defined by

$$\hat{\mathbf{Y}}_1 = \mathcal{H}_1(\hat{\mathbf{X}}_1). \quad (8)$$

that would be observed by the sensor \mathbf{S}_1 at time t_1 is expected to be practically indistinguishable from the observed image \mathbf{Y}_1 . This consistency property is expressed as

$$\hat{\mathbf{Y}}_1 \approx \mathbf{Y}_1. \quad (9)$$

It is the main rationale of the strategy introduced in [11], where CD is conducted by alternating between the CM inference and the fusion steps. These two steps are achieved by iteratively solving two optimization problems whose complexities depend

on the acquisition scenario. In the next section, we show that the robust-fusion based CD approach can be envisioned from an adversarial paradigm, by jointly conducting CM inference and fusion in a unified framework. In particular, this framework embeds two essential building blocks dedicated to the two aforementioned steps. Notably, since one of these tasks consists in solving a fusion problem, any off-the-shelf pre-trained network designed for this purpose can be effectively reused as an essential building block. This framework is detailed in what follows.

3. Proposed adversarial framework

The previous section showed that CD between images acquired by heterogeneous sensors can be conducted by inferring a latent image pair $\{\mathbf{X}_1, \mathbf{X}_2\}$, or one latent image and the CI $\{\mathbf{X}_1, \Delta\mathbf{X}\}$ with $\mathbf{X}_2 = \mathbf{X}_1 + \Delta\mathbf{X}$. This inference problem, coined as robust fusion, relies on two steps, namely CI inference and fusion, detailed in Sections 2.1 and 2.2, respectively. This paper shows that these two steps can be performed by two essential building blocks embedded into an unifying adversarial architecture. This architecture offers a high degree of modularity and thus opens the door to countless refinement opportunities and subsequent performance improvements. The different building blocks of the adversarial architecture, depicted in Fig. 1 and referred to as CD-GAN in the sequel, as well as the training strategy are detailed in this section. Finally, we provide a specific instance of this framework for the particular case of complementary acquisitions.

3.1. Overall architecture

CI inference – Estimating the CI $\Delta\mathbf{X}$ from the pair of observed images $(\mathbf{Y}_1, \mathbf{Y}_2)$ can be interpreted as a mapping from the product set $\mathbb{R}^{m_1 \times n_1} \times \mathbb{R}^{m_2 \times n_2}$ of observed images towards the set $\mathbb{R}^{n \times m}$ of CIs. In the proposed framework, this unknown mapping is defined as

$$\Delta\hat{\mathbf{X}} = \mathbf{C}(\mathbf{Y}_1, \mathbf{Y}_2; \boldsymbol{\Theta}_C) \quad (10)$$

where $\mathbf{C}(\cdot, \cdot; \boldsymbol{\Theta}_C)$ is a deep network parameterized by $\boldsymbol{\Theta}_C$. This network is able to solve the CD problem in the case of heterogeneous sensors. Indeed, once trained, it directly provides the estimated CI $\Delta\hat{\mathbf{X}}$, from which the binary change map $\hat{\mathbf{d}}$ can be estimated as detailed in Section 2.1. In a supervised setting, training this network would require triplets $(\mathbf{Y}_1^{(i)}, \mathbf{Y}_2^{(i)}, \Delta\mathbf{X}_1^{(i)})$ composed of a pair of observed images and the corresponding change image, which may significantly limit the applicability of the method. Conversely, the proposed adversarial approach will propose to train this network within an unsupervised setting, i.e., without requiring the availability of such triplets.

Fusion – Using the notations introduced in the previous section, the fusion step can be formulated as a mapping from the product set $\mathbb{R}^{m_1 \times n_1} \times \mathbb{R}^{m_2 \times n_2}$ of observed images to the set $\mathbb{R}^{n \times m}$ of latent images, i.e.,

$$\hat{\mathbf{X}}_1 = \mathbf{F}(\mathbf{Y}_1, \tilde{\mathbf{Y}}_2; \boldsymbol{\Theta}_F), \quad (11)$$

where $\mathbf{F}(\cdot, \cdot; \boldsymbol{\Theta}_F)$ stands for a fusion method parameterized by $\boldsymbol{\Theta}_F$. In this work, we propose to benefit from recent efforts

¹In this paper, we introduced the corrected image at time t_1 . Another possible choice would consist in defining the corrected image $\tilde{\mathbf{Y}}_1 \triangleq \mathbf{Y}_1 + \mathcal{H}_1(\Delta\mathbf{X})$ that would be acquired by the sensor \mathbf{S}_1 at time t_2 , i.e., after the changes occur. In this case, the subsequent technical derivations should be adapted accordingly.

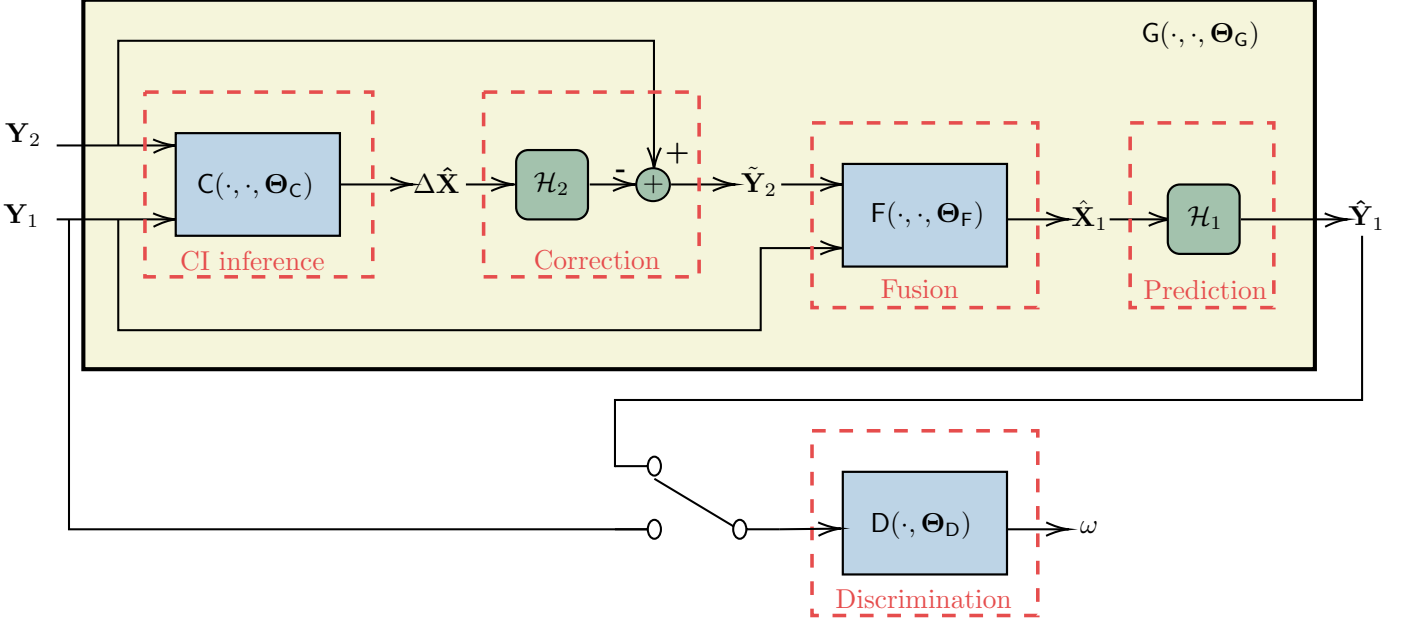


Figure 1: Flowchart of the overall proposed CD-GAN architecture. The 3 main networks $C(\cdot, \cdot; \Theta_C)$, $F(\cdot, \cdot; \Theta_F)$ and $D(\cdot, \cdot; \Theta_D)$ appear as blue light boxes. The virtual generator $G(\cdot, \cdot; \Theta_G)$ is figured as a beige box. The five main sub-tasks of the architecture are highlighted as red dashed-lined boxes.

promoting deep architectures specifically designed to solve the fusion task [50, 7, 49, 57]. More precisely, in the sequel of the paper, we will assume that the fusion method $F(\cdot, \cdot; \Theta_F)$ can be chosen as any off-the-shelf pre-trained deep network. Thus Θ_F remains fixed in the following. It is worth noting that this fusion network implicitly requires the knowledge of the CI ΔX (or one estimate $\Delta \hat{X}$), which is provided by the network $C(\cdot, \cdot; \Theta_C)$. Indeed, this network takes as one input the corrected image \tilde{Y}_2 which depends on the CI ΔX according to (7).

Fusion consistency-based training – A naive strategy to train the network $C(\cdot, \cdot; \Theta_C)$ dedicated to CI inference would consist in ensuring that the output \hat{X}_1 of the subsequent correction and fusion steps defined by (11) closely matches the true latent image X_1 . This approach, generally adopted by the GAN-based fusion methods from the literature, relies on an appropriate discriminator, denoted as $D(\cdot; \Theta_D)$ in Fig. 1. It would require the availability of a training set of triplets $(Y_1^{(i)}, Y_2^{(i)}, X_1^{(i)})$ composed of the two observed images and the associated latent image standing for the expected fusion result. Such a supervised learning appears restrictive since, in practice, it requires the availability of representative labelled data. This work proposes an alternative training that only requires the observed image Y_1 . More precisely, by leveraging the fusion consistency (9) discussed in Section 2.3, one designs a discriminative network which assesses the quality of the predicted image \hat{Y}_1 defined by (8) with respect to (w.r.t.) the observed image Y_1 . This can be achieved by designing a discriminative network $D(\cdot; \Theta_D)$ such that

$$\omega = D(Y; \Theta_D) \quad (12)$$

where $Y \in \{Y_1, \hat{Y}_1\}$, $\omega \in \{0, 1\}$ is the binary label stating the likeness between the predicted image \hat{Y}_1 and the observed

image Y_1 ($\omega = 1$ meaning “alike” and $\omega = 0$ meaning “different”), and Θ_D is the set of the discriminative network parameters.

To summarize, the overall proposed adversarial architecture can be decomposed into five main sub-tasks depicted as red dashed-lined boxes in Fig. 1 and listed below:

- *CI inference*: the network $C(\cdot, \cdot; \Theta_C)$ produces a candidate CI $\Delta \hat{X}$ from the two observed images Y_1 and Y_2 ,
- *Correction*: the inferred candidate CI $\Delta \hat{X}$ is used to produce the corrected image \tilde{Y}_2 following (7),
- *Fusion*: the corrected image \tilde{Y}_2 is subsequently fused with the observed image Y_1 thanks to the pre-trained network $F(\cdot, \cdot; \Theta_F)$, providing the estimated latent image \hat{X}_1 ,
- *Prediction*: the predicted image $\hat{Y}_1 = \mathcal{H}_1(\hat{X}_1)$ is derived,
- *Discrimination*: the quality of the predicted image \hat{Y}_1 w.r.t. the observed image Y_1 is assessed thanks to the discriminative network $D(\cdot; \Theta_D)$.

A meaningful interpretation of the proposed architecture can be drawn by observing that the sub-network highlighted as a beige box in Fig. 1 and denoted $G(\cdot, \cdot; \Theta_G)$ acts as a deterministic generator with parameters $\Theta_G = \{\Theta_C, \Theta_F\}$. It takes the two observed images Y_1 and Y_2 as inputs and embeds the four first sub-tasks of the architecture to produce the predicted image \hat{Y}_1 . This can be written as

$$G(Y_1, Y_2; \Theta_G) \triangleq \hat{Y}_1, \quad (13)$$

where we recall that

$$\begin{aligned}\hat{\mathbf{Y}}_1 &= \mathcal{H}(\hat{\mathbf{X}}_1) \\ \hat{\mathbf{X}}_1 &= \mathbf{F}(\mathbf{Y}_1, \hat{\mathbf{Y}}_2; \boldsymbol{\Theta}_F) \\ \hat{\mathbf{Y}}_2 &= \mathbf{Y}_2 - \mathcal{H}_2(\Delta\hat{\mathbf{X}}) \\ \Delta\hat{\mathbf{X}} &= \mathbf{C}(\mathbf{Y}_1, \mathbf{Y}_2; \boldsymbol{\Theta}_C)\end{aligned}$$

or, compactly written,

$$\mathbf{G}(\mathbf{Y}_1, \mathbf{Y}_2; \boldsymbol{\Theta}_G) = \mathcal{H}(\mathbf{F}(\mathbf{Y}_1, \mathbf{Y}_2 - \mathcal{H}_2(\mathbf{C}(\mathbf{Y}_1, \mathbf{Y}_2; \boldsymbol{\Theta}_C)); \boldsymbol{\Theta}_F)).$$

Since the fusion network is assumed to have been pre-trained, the sole generator parameters that still require training are the parameters $\boldsymbol{\Theta}_C$ of the CI inference network. Thus, in what follows, this generator will be denoted $\mathbf{G}(\cdot, \cdot; \boldsymbol{\Theta}_C)$. Introducing this generator allows the proposed architecture to be framed into an (almost) conventional GAN framework².

3.2. Loss function

As previously stated, the fusion network $\mathbf{F}(\cdot, \cdot; \boldsymbol{\Theta}_F)$ can be chosen as any state-of-the-art network (or even any model-based fusion algorithm) from the literature and it is assumed to have been pre-trained (or calibrated) beforehand. In this framework, two sub-networks need to be trained, namely the CI inference network $\mathbf{C}(\cdot, \cdot; \boldsymbol{\Theta}_C)$ and the discriminative network $\mathbf{D}(\cdot; \boldsymbol{\Theta}_D)$, by minimizing a well-chosen loss function. First of all, the proposed loss function incorporates an adversarial cost defined by

$$\begin{aligned}\mathcal{L}_{\text{adv}}(\boldsymbol{\Theta}_C, \boldsymbol{\Theta}_D) &= \mathbb{E}_{\mathbf{Y}_1} [\log \mathbf{D}(\mathbf{Y}_1; \boldsymbol{\Theta}_D)] \\ &+ \mathbb{E}_{\mathbf{Y}_1, \mathbf{Y}_2} [\log(1 - \mathbf{D}(\mathbf{G}(\mathbf{Y}_1, \mathbf{Y}_2; \boldsymbol{\Theta}_C); \boldsymbol{\Theta}_D))].\end{aligned}\quad (14)$$

Then, as emphasized by previous works [19], the adversarial cost can benefit from its combination with application-oriented costs. The first one is the so-called prediction loss introduced to assess the quality of the predicted image $\hat{\mathbf{Y}}_1$ w.r.t. the observed image \mathbf{Y}_1

$$\mathcal{L}_{\text{pre}}(\boldsymbol{\Theta}_C) = \mathbb{E}_{\mathbf{Y}_1, \mathbf{Y}_2} [\|\mathbf{Y}_1 - \hat{\mathbf{Y}}_1\|_F^2] \quad (15)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Since the changes are expected to affect only a few pixels, as discussed in Section 2.1, a second term promoting spatial sparsity of the estimated CI is also introduced [12, 10]

$$\mathcal{L}_{\text{spa}}(\boldsymbol{\Theta}_C) = \mathbb{E}_{\mathbf{Y}_1, \mathbf{Y}_2} [\|\Delta\hat{\mathbf{X}}\|_{2,1}] \quad (16)$$

where $\|\cdot\|_{2,1}$ denotes a group sparsity promoting norm. Finally, training both networks is formulated as the minimax problem

$$\min_{\boldsymbol{\Theta}_C} \max_{\boldsymbol{\Theta}_D} \mathcal{L}_{\text{adv}}(\boldsymbol{\Theta}_C, \boldsymbol{\Theta}_D) + \alpha \mathcal{L}_{\text{pre}}(\boldsymbol{\Theta}_C) + \beta \mathcal{L}_{\text{spa}}(\boldsymbol{\Theta}_C), \quad (17)$$

where α and β are hyperparameters adjusting the relative weights of the terms. In practice, following the common training procedures of GAN-based architectures, the two networks are alternately updated by stochastic gradient-based optimization, i.e., the expectations involved in (17) are approximated by empirical averages over minibatches.

Remark (Interpreting the generator training) Interestingly, the minimax optimization (17) underlying the generative-discriminative training can be interpreted from an image processing perspective. Indeed, let consider the deterministic setting or, equivalently, a minibatch composed of a unique sample. Under this simplifying assumption, training the generative network $\mathbf{G}(\cdot, \cdot; \boldsymbol{\Theta}_C)$ dedicated to CI inference only boils down to solving the optimization problem

$$\min_{\boldsymbol{\Theta}_C} \|\mathbf{Y}_1 - \hat{\mathbf{Y}}_1\|_F^2 + \|\Delta\hat{\mathbf{X}}\|_{2,1} + \mathcal{L}_{\text{adv}}(\boldsymbol{\Theta}_C), \quad (18)$$

where the predicted image $\hat{\mathbf{Y}}_1$ and the CI $\Delta\hat{\mathbf{X}}$ implicitly depend on $\boldsymbol{\Theta}_C$. More precisely, since $\hat{\mathbf{Y}}_1 = \mathcal{H}_1(\hat{\mathbf{X}}_1)$, the loss function can be rewritten as

$$\min_{\boldsymbol{\Theta}_C} \|\Delta\check{\mathbf{Y}}_1 - \mathcal{H}_1(\Delta\hat{\mathbf{X}})\|_F^2 + \|\Delta\hat{\mathbf{X}}\|_{2,1} + \mathcal{L}_{\text{adv}}(\boldsymbol{\Theta}_C), \quad (19)$$

where $\check{\mathbf{Y}}_1 \triangleq \mathcal{H}_1(\mathbf{X}_2)$ is the image that would be observed by the sensor \mathcal{S}_1 at time t_2 and $\Delta\check{\mathbf{Y}}_1 \triangleq \check{\mathbf{Y}}_1 - \mathbf{Y}_1$ is the \mathcal{S}_1 -virtual CI (i.e., the difference between the image actually observed at time t_1 by the sensor \mathcal{S}_1 and the image virtually observed at time t_2 by the same sensor). Since the CI $\Delta\hat{\mathbf{X}}$ is fully defined as the output of the network $\mathbf{C}(\cdot, \cdot; \boldsymbol{\Theta}_C)$ according to (10), training this network through the nonlinear optimization problem (19) can be interpreted as a parametric inversion task: given the \mathcal{S}_1 -virtual CI $\Delta\check{\mathbf{Y}}_1$, one aims at inverting the measurement operator $\mathcal{H}_1(\cdot)$ seeking the solution in a parametric form specified by $\boldsymbol{\Theta}_C$. It is worth noting that this inversion step has been explicitly referred to as a correction in [12, 10]. Note that in (19), we have introduced an additional task-driven data fitting term: the adversarial term $\mathcal{L}_{\text{adv}}(\boldsymbol{\Theta}_C)$.

3.3. Detailed architecture in case of complementary acquisitions

This section provides details on the network architectures for a specific instance of the proposed framework. More precisely, it considers the particular case of complementary acquisitions, i.e., when the observed images \mathbf{Y}_1 and \mathbf{Y}_2 are LRHS and HRLS images respectively. As stated before, the fusion network $\mathbf{F}(\cdot, \cdot; \boldsymbol{\Theta}_F)$ can be chosen as any state-of-the-art architecture (or model-based algorithm) [27, 50, 7, 54, 56, 49, 17] and is assumed to have been trained (or calibrated) beforehand. In this work, as an example but without loss of generality, we adopt a network similar to the one introduced in [27], whose detailed architecture is depicted in Fig. 2. Specifically, the LRHS and HRLS images are provided as inputs of two respective sub-networks designed for feature extraction. Each sub-network consists of two successive convolutional layers followed by a leaky rectified linear unit (LeakyReLU) [30] and a down-convolution layer. Since the spatial resolution of the LRHS image \mathbf{Y}_1 is lower than the one of the target latent image \mathbf{X}_1 , the second convolution layer of the dedicated sub-network is an up-convolution layer. The two resulting feature maps are then concatenated and pass through a series of convolutional layers, along with skip connections [38]. Finally, ReLU is applied in the last layer to ensure nonnegativity of the output fused image.

²Another possible interpretation consists in pipelining the four last sub-tasks of the architecture to define an extended discriminator combined with a generator defined as the sole network dedicated to the CI inference

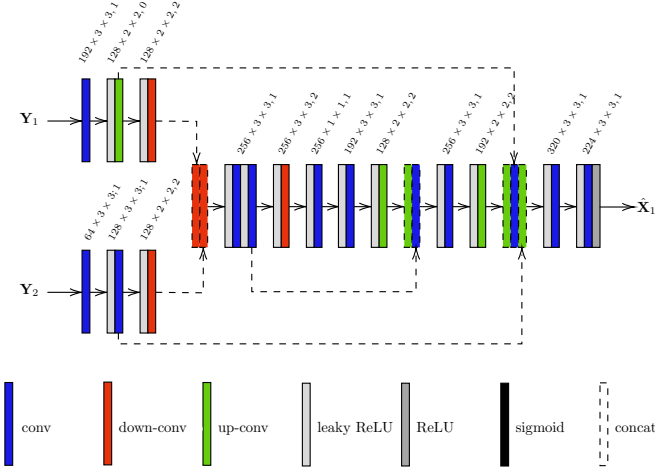


Figure 2: Architecture of the network $F(\cdot, \cdot; \Theta_F)$ dedicated to the fusion sub-task. The architecture of the network $C(\cdot, \cdot; \Theta_C)$ dedicated to CI inference only differs by the last ReLU layer that has been removed.

Regarding the network $C(\cdot, \cdot; \Theta_C)$ dedicated to CI inference, it basically performs a mapping with the same input and target spaces as the fusion network detailed above. Even if the respective underlying tasks are different (i.e., CD vs. fusion), their objectives are similar, e.g., extracting relevant spatial-spectral information from a pair of images with different spatial and spectral resolutions to produce a HRHS (latent or change) image. Thus, it is reasonable to adopt a similar architecture for $C(\cdot, \cdot; \Theta_C)$ as for $F(\cdot, \cdot; \Theta_F)$. The only difference lies in the last layer: the ReLU has been removed, as the CI is not necessarily nonnegative.

Finally, the discriminative network $D(\cdot; \Theta_D)$ takes an LRHS image as an input, the observed image Y_1 or the predicted image \hat{Y}_1 , to provide a binary decision $\omega \in \{0, 1\}$. When $\omega = 1$, the network decides the input image is an actually observed one. Conversely, the discriminator provides the output $\omega = 0$ when it decides the input image is a predicted image. As depicted in Fig. 3, it consists of three down-convolution layers and two flat convolution layers. The last layer of the network includes a sigmoid activation function.




Figure 3: Architecture of the discriminative network.

4. Experiments on simulated data sets

Statistical performance assessment of newly developed methods requires ground truth data to compute quantitative figures-of-merit. In the context of CD, ground truth data should provide the location of the changes that actually occurred in a

geographical area between the two image acquisitions. However, in practice, it is generally difficult to obtain enough pairs of real images acquired by heterogeneous sensors over the same geographical area at different times. Moreover, it is even more difficult to obtain the associated ground truth binary maps locating the changes between the acquisitions. As an alternative, this section aims at assessing the efficiency of the proposed CD-GAN framework thanks to experiments conducted on synthetic data. It compares its performance with standard and state-of-the-art CD techniques. In this perspective, the proposed framework is instantiated for the particular case of complementary acquisitions. More precisely, throughout this section, the LRHS observed image Y_1 is an HS image of low spatial resolution, while the HRLS observed image Y_2 is an MS image of higher spatial resolution.

4.1. Synthetic data generation

The experiments reported hereafter have been conducted on synthetic data obtained following the simulation protocol proposed in [11]. This protocol relies on the availability of a single HRHS image X_{ref} to generate pairs of observed³ HRLS and LRHS images through an unmixing-upmixing process. This process introduces physically-inspired, realistic changes to the observed images, with predefined change maps that will be subsequently used to evaluate the performance of CD methods. The successive steps of this protocol are briefly sketched below. Interested readers are invited to refer to the work in [11] for more details.




Figure 4: Color composition of the three hyperspectral reference images.

Generation of HRHS images – Three sets of 22 HRHS reference images have been extracted from three real HS images acquired by the AVIRIS sensor, composed of 789×789 pixels and 224 spectral bands. They are depicted in Fig. 4. Each reference image denoted by X_{ref} is composed of $m = 224$ spectral bands and are of size 120×120 pixels (i.e., $n = 14400$). The resulting 66 reference images are then processed as follows.

Unmixing – From each reference image X_{ref} , a set of k endmembers gathered in an $m \times k$ matrix M_{ref} are extracted using vertex component analysis [33], where the number k of endmembers has been adjusted using Hysime [32]. Based on this set of endmembers, the reference image has been unmixed

³To lighten the writing, we took the licence to refer to the synthetically generated HRLS and LRHS images as “observed” or “acquired” images.

using SUnSAL [18] to recover a corresponding $k \times n$ abundance matrix \mathbf{A}_{ref} .

Generation of changes – Three kinds of simulated yet realistic changes are independently applied to the reference abundance matrices \mathbf{A}_{ref} to produce modified abundance matrices denoted \mathbf{A}_{chg} . These three change rules are referred to as *zero abundance* (R_z), *same abundance* (R_s) and *block abundance* (R_b) in the sequel (see [11] for more details). For each reference image, the pixels whose respective abundance vectors have been modified are identified by non-zero values in the predefined reference binary map $\mathbf{d}_{\text{ref}} \in \{0, 1\}^n$.

Upmixing – Pairs $(\mathbf{X}_1, \mathbf{X}_2)$ of HRHS latent images are computed by linearly mixing endmember matrices \mathbf{M}_{ref} with the abundances in \mathbf{A}_{ref} or \mathbf{A}_{chg} . More precisely, these simulated pairs of latent images \mathbf{X}_1 and \mathbf{X}_2 are defined as

$$(\mathbf{X}_1, \mathbf{X}_2) \in \left\{ (\mathbf{M}_{\text{ref}}\mathbf{A}_{\text{ref}}, \mathbf{M}_{\text{ref}}\mathbf{A}_{\text{chg}}), (\mathbf{M}_{\text{ref}}\mathbf{A}_{\text{chg}}, \mathbf{M}_{\text{ref}}\mathbf{A}_{\text{ref}}) \right\}. \quad (20)$$

Generation of observed images – Given the HRHS latent images \mathbf{X}_1 and \mathbf{X}_2 produced as above, respective pairs $(\mathbf{Y}_1, \mathbf{Y}_2)$ of observed images are generated according to the forward models (2). In these experiments, as stated above, the observed images are assumed to be of complementary resolutions. Thus, the spatial degradation operator $\mathcal{H}_1(\cdot)$ is chosen as a spatially-invariant Gaussian blur with standard deviation $\sigma = 2.35$ followed by a regular down-sampling in both directions of factor 4 leading to $n_2 = 16n_1$ (see (1)). The spectral degradation operator $\mathcal{H}_2(\cdot)$ mimics the response of a MS sensor by averaging four contiguous bands along the spectral dimension.

Finally, with the above steps, 396 pairs of images have been generated, from which 376 pairs have been used for training, and the other for testing.

4.2. Experimental settings

The proposed CD-GAN architecture is implemented using the PyTorch framework on a computer equipped with a Quadro RTX 6000 GPU. Unless otherwise stated, the hyperparameters adjusting the terms associated with the prediction and the spatial regularizations in the loss function (17) have been set to $\alpha = 1$ and $\beta = 10^{-3}$. During the training, Adam [21] is considered as the optimizer with an initial learning rate set to 2×10^{-4} over 15 epochs and batches composed of 4 training samples. These parameters have been adopted for all experiments reported in what follows. The change maps estimated by the proposed method is denoted as $\hat{\mathbf{d}}_{\text{CD-GAN}}$.

Compared methods – The proposed method is compared to five other ones. However, very few works considered the problem of CD between images of different resolutions. Most methods apply independent preprocessing to the two observed images to reach the same spatial and the same spectral resolutions which makes a pixel-wise comparison possible.

- The first method is the fusion-based approach proposed in [11]. It first fuses the two observed images using the fast algorithm based on the solving of the Sylvester equation (FUSE) proposed in [52] to infer a common latent HRHS image. Then, an LRHS change image is computed from the spatially degraded latent image and the observed LRHS image \mathbf{Y}_1 . Finally, a spatially regularized CVA (sCVA) [20] has been conducted to produce an estimated LR CM denoted $\hat{\mathbf{d}}_{\text{F-FUSE}}$.
- The second method also elaborates on the first one but it embeds the dual spatial-spectral pyramid network With transformer for hyperspectral image fusion (DSP-NET) [57]. The estimated LR CM is denoted $\hat{\mathbf{d}}_{\text{F-SSRNET}}$.
- The third method also elaborates on the first one but it embeds the dual spatial-spectral pyramid network with transformer for hyperspectral image fusion (DSP-NET) [44]. The estimated LR CM is denoted $\hat{\mathbf{d}}_{\text{F-DSPNET}}$.
- The fourth method is the image translation-based unsupervised change detection proposed in [28]. The generated change map is denoted as $\hat{\mathbf{d}}_{\text{USCD}}$.
- The fifth method applies the same operations as the second one but in a reverse order. The observed image \mathbf{Y}_1 is first spectrally degraded by applying $\mathcal{H}_2(\cdot)$ and then spatially superresolved using [58]. The pair composed of the resulting HRHS image and the observed image \mathbf{Y}_2 is then analyzed using sCVA to produce a HR change map denoted $\hat{\mathbf{d}}_{\text{LSHR}}$.

Figures-of-merit – To quantitatively evaluate the performance of the CD methods, the estimated change maps $\hat{\mathbf{d}}$ are compared to the reference change map \mathbf{d}_{ref} . This allows us to derive the empirical receiver operating characteristics (ROC) curves. These ROC curves represent the estimated probability of detection P_D as a function of the estimated probability of false alarm P_{FA} . Detection occurs when an actually changed pixel is identified as changed, while false alarm occurs when an unchanged pixel is identified as changed. Empirical ROC curves are the privileged figure-of-merit for detection performance assessment [37]. Contrary to classification-oriented metrics, a ROC curve comprehensively displays the estimated detection performance in terms of possible trade-offs between detection and false alarms on a single plot. Considering a wide range of threshold values allows to cover these possible trade-offs. Two quantitative metrics are also derived from these ROC curves, namely the area under the curve (AUC), sometimes referred to as the c-statistic [15, Chap. 9], and the distance (dist.) between the no detection point and the point at the interception of the ROC curve. For both metrics, the closer the values are to 1, the better the CD methods. In addition, the performances of the compared methods are evaluated w.r.t. the overall accuracy (OA) defined as

$$\text{OA} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (21)$$

where TP and TN refer to the true positive and true negative, i.e., the numbers of correctly identified changed and unchanged pixels, respectively, while FP and FN refer to the false positive and false negative, i.e., the numbers of incorrectly identified unchanged and changed pixels, respectively.

4.3. Results

Performance comparison – The ROC curves obtained by the compared methods are shown in Fig. 5, for the three change rules R_b , R_s and R_z . The associated quantitative results are reported in Table 1. Clearly, these first results show the superiority of the proposed CD-GAN framework when compared to the five other CD methods. More precisely, CD-GAN provides better detection even for very low probability of false alarm with an acceptable time. The method based on \hat{d}_{LSHR} is faster than other methods, behave similarly and under-perform the other methods. Fig. 6 shows the true CM and the CM estimated by the compared methods. Note that the estimated CMs \hat{d}_{F-FUSE} , $\hat{d}_{F-SSRNET}$ are defined at low spatial resolutions, contrary to the other ones.

| | | \hat{d}_{CD-GAN} | \hat{d}_{F-FUSE} | $\hat{d}_{F-SSRNET}$ | $\hat{d}_{F-DSPNET}$ | \hat{d}_{USCD} | \hat{d}_{LSHR} |
|---------|-------|--------------------|--------------------|----------------------|----------------------|------------------|------------------|
| R_b | AUC | 0.9962 | 0.9660 | 0.9342 | 0.9596 | 0.6906 | 0.8081 |
| | Dist. | 0.9774 | 0.9190 | 0.9286 | 0.8955 | 0.6756 | 0.7248 |
| | OA | 0.9879 | 0.9806 | 0.9809 | 0.9830 | 0.9781 | 0.9768 |
| R_s | AUC | 0.9999 | 0.9971 | 0.9968 | 0.9969 | 0.9626 | 0.9965 |
| | Dist. | 0.9988 | 0.9914 | 0.9914 | 0.9841 | 0.9626 | 0.9841 |
| | OA | 0.9963 | 0.9869 | 0.9869 | 0.9938 | 0.9803 | 0.9898 |
| R_z | AUC | 0.9997 | 0.9956 | 0.9957 | 0.9844 | 0.9930 | 0.9843 |
| | Dist. | 0.9941 | 0.9864 | 0.9957 | 0.9612 | 0.9930 | 0.9612 |
| | OA | 0.9972 | 0.9861 | 0.9859 | 0.9955 | 0.9902 | 0.9780 |
| Time(s) | | 1.9079 | 0.3899 | 0.2414 | 0.7005 | 0.1667 | 0.0099 |

Table 1: Quantitative detection performance and computational times.

Sensitivity analysis – We now discuss the impact of the three terms in the loss function in (17), namely, the adversarial loss $\mathcal{L}_{adv}(\cdot, \cdot)$, the prediction loss $\mathcal{L}_{pre}(\cdot)$ and the spatial sparsity regularization $\mathcal{L}_{spa}(\cdot)$. Table 2 reports the quantitative results associated with the three change rules R_z , R_s and R_b while adopting different loss terms in (17): first column without the adversarial loss, second column without the prediction loss and third one without spatial sparsity regularization. Optimal values are highlighted in bold. Clearly, in this case, the prediction loss and the spatial sparsity regularization play important roles in the proposed CD-GAN, whereas combining the three terms leads to the best results.

In addition, we also analyze in this table the impact of the spatial sparsity regularization in the loss function (17) by adjusting the hyperparameter β . The empirical results demonstrate the interest of incorporating this regularization since the case $\beta = 0$, i.e., with no regularization, leads to the worst detection performance. Conversely, for a quite wide range of

non-zero values of β , i.e., $\beta \in \{10^{-4}, 10^{-3}, 10^{-2}\}$, the detection performance is clearly better with the best one obtained for $\beta = 10^{-3}$, which is the default value chosen for the experiments reported in this section.

| | | no adv. | $\alpha = 0$ | $\beta = 0$ | $\beta = 10^{-4}$ | $\beta = 10^{-3}$ | $\beta = 10^{-2}$ |
|-------|-------|---------|--------------|-------------|-------------------|-------------------|-------------------|
| R_b | AUC | 0.9954 | 0.9324 | 0.9335 | 0.9925 | 0.9962 | 0.9950 |
| | Dist. | 0.9762 | 0.9102 | 0.8357 | 0.9688 | 0.9774 | 0.9651 |
| | OA | 0.9858 | 0.9804 | 0.9805 | 0.9836 | 0.9879 | 0.9853 |
| R_s | AUC | 0.9910 | 0.9158 | 0.8399 | 0.9835 | 0.9999 | 0.9896 |
| | Dist. | 0.9718 | 0.9031 | 0.7892 | 0.9601 | 0.9988 | 0.9560 |
| | OA | 0.9862 | 0.9816 | 0.9765 | 0.9798 | 0.9963 | 0.9786 |
| R_z | AUC | 0.9949 | 0.9273 | 0.8854 | 0.9768 | 0.9997 | 0.9937 |
| | Dist. | 0.9724 | 0.9165 | 0.8534 | 0.9350 | 0.9941 | 0.9684 |
| | OA | 0.9907 | 0.9836 | 0.9829 | 0.9818 | 0.9972 | 0.9932 |

Table 2: Sensitivity analysis w.r.t. the hyperparameters in CD-GAN: quantitative detection performance.

Note that a similar analysis has been conducted for the hyperparameter α adjusting the weight of the prediction term in the overall loss (17). The results, which show the limited impact of this hyperparameter on the detection performance, are not reported herein for brevity.

Impact of the fusion methods – In the proposed CD-GAN framework, the fusion network $F(\cdot, \cdot; \Theta_F)$ can be chosen by the end-user and has been assumed to be trained beforehand. Section 3.3 describes the fusion network adopted in most of the experiments discussed herein. This network heavily relies on the PS-GAN architecture proposed in [27]. To illustrate its versatility, we propose to instantiate the proposed CD framework with three other fusion networks, namely the deep blind hyperspectral image fusion (DBIN) proposed in [50], the SSR-NET network proposed in [57] and the DSP-NET network proposed in [44]. Besides, as a complementary analysis, we consider a semi-supervised scenario where the latent image \mathbf{X}_1 (i.e., the HRHS image associated with the observed image \mathbf{Y}_1) is also available for training. Note that this more favorable scenario is the one considered in [27] to train PS-GAN. In this case, the discriminative network $D(\cdot; \Theta_D)$ and the prediction loss (15) can be easily adapted accordingly to distinguish the estimated fused image $\hat{\mathbf{X}}_1$ from the true latent image \mathbf{X}_1 . The quantitative metrics are reported in Table 3 for the three implemented change rules. These results show that, whatever the adopted fusion network, the proposed unsupervised CD-GAN framework reaches detection performance comparable to the one obtained in a more favorable semi-supervised scenario. In addition, they show that the proposed CD framework is quite robust to the choice of the fusion method.

Robustness w.r.t. the forward models – The proposed framework sketched in Fig. 1 requires the definition of the spatial and spectral degradation operators $\mathcal{H}_1(\cdot)$ and $\mathcal{H}_2(\cdot)$. In some practical applications, one may have only a partial knowledge




Figure 5: ROC curves obtained by the compared CD methods for the three change rules: R_b (left), R_s (middle) and R_z (right).




Figure 6: CM estimated by the compared methods on the data sets generated according to the three change rules, namely R_b (1st row), R_s (2nd row) and R_z (3rd row) with from left to right: ground truth, \hat{d}_{CD-GAN} , \hat{d}_{F-FUSE} , $\hat{d}_{F-SSRNET}$, $\hat{d}_{F-DSPNET}$, \hat{d}_{USCD} , \hat{d}_{LRHS} .

| | | Unsupervised CD-GAN | | | | Semi-supervised CD-GAN |
|-------|-------|---------------------|---------------|---------|---------------|---------------------------|
| | | DBIN | SSR-NET | DSP-NET | PS-GAN | |
| R_b | AUC | 0.9860 | 0.9962 | 0.9948 | 0.9962 | 0.9971 |
| | Dist. | 0.9650 | 0.9683 | 0.9753 | 0.9774 | 0.9774 |
| | OA | 0.9753 | 0.9834 | 0.9874 | 0.9879 | 0.9832 |
| R_s | AUC | 0.9933 | 0.9974 | 0.9988 | 0.9999 | 0.9946 |
| | Dist. | 0.9750 | 0.9675 | 0.9938 | 0.9988 | 0.9737 |
| | OA | 0.9867 | 0.9936 | 0.9929 | 0.9967 | 0.9902 |
| R_z | AUC | 0.9975 | 0.9984 | 0.9986 | 0.9997 | 0.9940 |
| | Dist. | 0.9849 | 0.9954 | 0.9883 | 0.9941 | 0.9851 |
| | OA | 0.9927 | 0.9936 | 0.9939 | 0.9961 | 0.9902 |

Table 3: Impact of the fusion method: quantitative detection performance.

have been conducted to assess the robustness of the proposed CD-GAN framework w.r.t. model mismatch. More precisely, first the CD-GAN architecture has been trained with corrupted degradation operators on a training set composed of observed images generated with uncorrupted degradation operators and following the protocol described in Section 4.1. Then, with all network parameters fixed, the corrupted operators have been updated in a second training procedure. Finally, the CD-GAN has been trained a second time based on these corrected operators. Regarding the corruption processes, the spatial degradation operator $\mathcal{H}_1(\cdot)$ specified in Section 4.1 has been replaced by a spatially-invariant Gaussian blur with standard deviation $\sigma = 1.70$. Inspired by the study conducted in [51], the spectral filters which compose the spectral degradation operator $\mathcal{H}_2(\cdot)$ have been corrupted by an additive zero-mean Gaussian noise with a variance leading to a SNR = 8dB.

Table 4 reports the quantitative results obtained by the compared methods. They show that the proposed CD-GAN obtains the highest metrics. In conclusion, when compared with other

regarding these operators related to the sensors. Experiments

| | | $\hat{\mathbf{d}}_{\text{CD-GAN}}$ | $\hat{\mathbf{d}}_{\text{F-FUSE}}$ | $\hat{\mathbf{d}}_{\text{F-SSRNET}}$ | $\hat{\mathbf{d}}_{\text{F-DSPNET}}$ | $\hat{\mathbf{d}}_{\text{USCD}}$ | $\hat{\mathbf{d}}_{\text{LSHR}}$ |
|-------|-------|------------------------------------|------------------------------------|--------------------------------------|--------------------------------------|----------------------------------|----------------------------------|
| R_b | AUC | 0.9948 | 0.7236 | 0.8638 | 0.8935 | 0.4655 | 0.5619 |
| | Dist. | 0.9714 | 0.6423 | 0.7858 | 0.7649 | 0.4635 | 0.4862 |
| | OA | 0.9914 | 0.9443 | 0.9713 | 0.9537 | 0.9142 | 0.9196 |
| R_s | AUC | 0.9918 | 0.9825 | 0.9847 | 0.9891 | 0.9421 | 0.9796 |
| | Dist. | 0.9876 | 0.9832 | 0.9734 | 0.9811 | 0.9349 | 0.9557 |
| | OA | 0.9900 | 0.9812 | 0.9823 | 0.9858 | 0.9727 | 0.9762 |
| R_z | AUC | 0.9934 | 0.9753 | 0.9811 | 0.9917 | 0.8017 | 0.7498 |
| | Dist. | 0.9812 | 0.9216 | 0.9653 | 0.9789 | 0.7983 | 0.6420 |
| | OA | 0.9914 | 0.9623 | 0.9838 | 0.9879 | 0.9705 | 0.9647 |

Table 4: Robustness w.r.t. forward models: quantitative detection performance.

methods, the proposed CD-GAN framework is shown to be robust w.r.t. the misspecification of the degradation operators for the three implemented change rules.

5. Experiments on real data sets

For an illustrative purpose, complementary experiments have been conducted on two real data sets, namely the Santa Barbara data set and Yancheng data set. Santa Barbara data set is composed of two HS images acquired by the AVIRIS sensor in 2013 and 2014, respectively, to monitor the change of composition over Santa Barbara area in California [40]. These images are composed of 984×740 pixels, with $m = 224$ spectral bands. In the experiments, two HRHS sub-images of sizes 120×120 are selected. Yancheng data set is composed of two HS images that have been captured on May 3, 2006 and April 23, 2007 by the Hyperion sensor. They reflect changes in farmland near the city of Yancheng, Jiangsu province, China. The spatial and spectral degradation operators $\mathcal{H}_1(\cdot)$ and $\mathcal{H}_2(\cdot)$ detailed in Section 4.1 are applied to each of these HRHS images to mimic heterogeneous acquisitions with complementary resolutions, i.e., producing one LRHS image and one HRLS image. The HRLS and LRHS images as well as the actual binary CMs constituting the three data sets, denoted SB1, SB2 and YC, are shown in Fig. 7. The binary CMs are designed as in [40].

Regarding these three data sets, the resulting binary CMs recovered by the compared methods are shown in Fig. 8. Visual inspection allows one to state that recovering the changed areas between the two observed images is a challenging task for all compared methods. The estimated CMs $\hat{\mathbf{d}}_{\text{F-DSPNET}}$, $\hat{\mathbf{d}}_{\text{USCD}}$ and $\hat{\mathbf{d}}_{\text{LSHR}}$ contain a lot of false alarms, i.e., pixels decided as changed which actually are not affected by any change. The other compared methods exhibit significantly less false alarms, in particular, the estimated CM $\hat{\mathbf{d}}_{\text{F-SSRNET}}$. However it is worth noting that the change maps $\hat{\mathbf{d}}_{\text{F-FUSE}}$ and $\hat{\mathbf{d}}_{\text{F-SSRNET}}$ are defined at low spatial resolution. The results associated with the proposed CD-GAN show accurate detection of the changes with a low number of false alarms. These findings are confirmed by the ROC curves depicted in Fig. 9 and the quantitative results reported in Table 5. From Fig. 9, it




Figure 7: Data sets SB1 (top), SB2 (middle), YC (bottom): color composition of the HRLS observed image (left), color composition of the LRHS observed image (middle) and the binary CM (right) where white (resp. black) pixels correspond to changed (resp. unchanged) areas.

appears that the proposed CD-GAN provides high detection rate whatever the functioning point of the detector (i.e., for all values of PFA). From a quantitative point-of-view, the proposed CD-GAN framework obtains higher dist. score, AUC and OA.




Figure 8: Data sets SB1 (top), SB2 (middle) and YC (bottom): CM estimated by the compared methods with, from left to right: ground truth, $\hat{\mathbf{d}}_{\text{CD-GAN}}$, $\hat{\mathbf{d}}_{\text{F-FUSE}}$, $\hat{\mathbf{d}}_{\text{F-SSRNET}}$, $\hat{\mathbf{d}}_{\text{F-DSPNET}}$, $\hat{\mathbf{d}}_{\text{USCD}}$, $\hat{\mathbf{d}}_{\text{LSHR}}$.

6. Conclusion

This paper introduced a robust fusion-based adversarial framework for detecting changes between heterogeneous images in an unsupervised scenario. The proposed approach capitalized on the availability of a predefined and previously trained fusion network. Following a robust fusion-based change detection strategy, this network was complemented with another network having the same architecture and specifically dedicated




Figure 9: Data sets SB1 (top left), SB2 (top right) and YC (bottom): ROC curves obtained by the compared methods.

| | | \hat{d}_{CD-GAN} | \hat{d}_{F-FUSE} | $\hat{d}_{F-SSRNET}$ | $\hat{d}_{F-DSPNET}$ | \hat{d}_{USCD} | \hat{d}_{LRHS} |
|-----|-------|--------------------|--------------------|----------------------|----------------------|------------------|------------------|
| SB1 | AUC | 0.9399 | 0.8263 | 0.8753 | 0.9033 | 0.6519 | 0.9033 |
| | Dist. | 0.8568 | 0.7124 | 0.8753 | 0.8091 | 0.6519 | 0.8091 |
| | OA | 0.9508 | 0.9093 | 0.9221 | 0.9329 | 0.8888 | 0.9222 |
| SB2 | AUC | 0.9577 | 0.9075 | 0.9370 | 0.9496 | 0.7199 | 0.9498 |
| | Dist. | 0.9057 | 0.7939 | 0.8623 | 0.8954 | 0.7153 | 0.8923 |
| | OA | 0.9715 | 0.9544 | 0.9513 | 0.9632 | 0.9374 | 0.9595 |
| YC | AUC | 0.9938 | 0.9843 | 0.9417 | 0.9846 | 0.8163 | 0.9683 |
| | Dist. | 0.9644 | 0.9475 | 0.9306 | 0.9221 | 0.7884 | 0.8849 |
| | OA | 0.9888 | 0.9875 | 0.9875 | 0.9804 | 0.9588 | 0.9747 |

Table 5: Data sets SB1, SB2 and YC: quantitative detection performance.

to detect change between images of possibly different spatial and spectral resolutions. The overall architecture was trained within an adversarial paradigm, enriching the canonical adversarial loss function with task-driven terms. Experiments conducted on simulated and real data sets illustrated the efficiency, the versatility and the robustness of the proposed change detection framework.

Acknowledgements

The authors would like to thank Florentin Coeurdoux (IRIT, University of Toulouse) for his relevant advice regarding the practical implementation of the proposed CD-GAN architecture.

References

[1] Alberga, V., Idrissa, M., Lacroix, V., Inglada, J., 2007a. Comparison of similarity measures of multi-sensor images for change detection applica-

tions, in: Proc. IEEE Int. Conf. Geosci. Remote Sens. (IGARSS), IEEE, pp. 2358–2361.

[2] Alberga, V., Idrissa, M., Lacroix, V., Inglada, J., 2007b. Performance estimation of similarity measures of multi-sensor images for change detection applications, in: Proc. IEEE Int. Workshop Analysis Multitemporal Remote Sensing Images (MultiTemp), IEEE, pp. 1–5.

[3] Bovolo, F., Bruzzone, L., 2006. A theoretical framework for unsupervised change detection based on change vector analysis in the polar domain. *IEEE Trans. Geosci. Remote Sens.* 45, 218–236.

[4] Bruzzone, L., Bovolo, F., 2012. A novel framework for the design of change-detection systems for very-high-resolution remote sensing images. *Proc. IEEE* 101, 609–630.

[5] Bruzzone, L., Prieto, D.F., 2002. An adaptive semiparametric and context-based approach to unsupervised change detection in multitemporal remote-sensing images. *IEEE Trans. Image Process.* 11, 452–466.

[6] Dian, R., Guo, A., Li, S., 2023. Zero-shot hyperspectral sharpening. *IEEE Trans. Patt. Anal. Mach. Intell.*

[7] Dian, R., Li, S., Kang, X., 2020. Regularizing hyperspectral and multispectral image fusion by CNN denoiser. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 1124–1135.

[8] Du, B., Ru, L., Wu, C., Zhang, L., 2019. Unsupervised deep slow feature analysis for change detection in multi-temporal remote sensing images. *IEEE Trans. Geosci. Remote Sens.* 57, 9976–9992.

[9] Du, P., Wang, X., Chen, D., Liu, S., Lin, C., Meng, Y., 2020. An improved change detection approach using tri-temporal logic-verified change vector analysis. *ISPRS J. Photogramm. Remote Sens.* 161, 278–293.

[10] Ferraris, V., Dobigeon, N., Chabert, M., 2020. Robust fusion algorithms for unsupervised change detection between multi-band optical images – A comprehensive case study. *Information Fusion* 64, 293–317.

[11] 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.* 56, 1566–1578.

[12] Ferraris, V., Dobigeon, N., Wei, Q., Chabert, M., 2017b. Robust fusion of multi-band images with different spatial and spectral resolutions for change detection. *IEEE Trans. Comput. Imag.* 3, 175–186.

[13] Gao, F., Dong, J., Li, B., Xu, Q., 2016. Automatic change detection in synthetic aperture radar images based on PCANet. *IEEE Geosci. Remote Sens. Lett.* 13, 1792–1796.

[14] Gueguen, L., Hamid, R., 2016. Toward a generalizable image representation for large-scale change detection: Application to generic damage analysis. *IEEE Trans. Geosci. Remote Sens.* 54, 3378–3387.

[15] Hastie, T., Tibshirani, R., Friedman, J.H., 2009. The elements of statistical learning: data mining, inference, and prediction. 2 ed., Springer, New York.

[16] Hegazy, I.R., Kaloop, M.R., 2015. Monitoring urban growth and land use change detection with GIS and remote sensing techniques in Daqahlia governorate Egypt. *Int. J. Sustainable Built Environment* 4, 117–124.

[17] Hu, J.F., Huang, T.Z., Deng, L.J., Jiang, T.X., Vivone, G., Chanussot, J., 2021. Hyperspectral image super-resolution via deep spatio-spectral attention convolutional neural networks. *IEEE Trans. Neural Netw. Learn. Syst.*

[18] Iordache, M.D., Bioucas-Dias, J.M., Plaza, A., 2011. Sparse unmixing of hyperspectral data. *IEEE Trans. Geosci. Remote Sens.* 49, 2014–2039.

[19] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A., 2017. Image-to-image translation with conditional adversarial networks, in: Proc. Int. Conf. on Computer Vision and Pattern Recognition (CVPR), pp. 1125–1134.

[20] Johnson, R.D., Kasischke, E., 1998. Change vector analysis: A technique for the multispectral monitoring of land cover and condition. *Int. J. Remote Sensing* 19, 411–426.

[21] Kingma, D., Ba, J., 2014. Adam: A method for stochastic optimization. *Proc. IEEE Int. Conf. Learn. Represent. (ICLR)*.

[22] Li, H., Gong, M., Zhang, M., Wu, Y., 2021. Spatially self-paced convolutional networks for change detection in heterogeneous images. *IEEE J. Sel. Topics Appl. Earth Observations Remote Sens.* 14, 4966–4979.

[23] Li, S., Dian, R., Liu, H., 2023. Learning the external and internal priors for multispectral and hyperspectral image fusion. *Science China Information Sciences* 66, 140303.

[24] Li, Z., Shi, W., Zhang, H., Hao, M., 2017. Change detection based on gabor wavelet features for very high resolution remote sensing images. *IEEE Geosci. Remote Sens. Lett.* 14, 783–787.

- [25] Liu, J., Gong, M., Qin, A.K., Tan, K.C., 2019. Bipartite differential neural network for unsupervised image change detection. *IEEE Trans. Neural Netw. Learn. Syst.* 31, 876–890.
- [26] Liu, J., Gong, M., Qin, K., Zhang, P., 2016. A deep convolutional coupling network for change detection based on heterogeneous optical and radar images. *IEEE Trans. Neural Netw. Learn. Syst.* 29, 545–559.
- [27] Liu, Q., Zhou, H., Xu, Q., Liu, X., Wang, Y., 2020. Psgan: A generative adversarial network for remote sensing image pan-sharpening. *IEEE Trans. Geosci. Remote Sens.* .
- [28] Liu, Z.G., Zhang, Z.W., Pan, Q., Ning, L.B., 2021. Unsupervised change detection from heterogeneous data based on image translation. *IEEE Trans. Geosci. Remote Sens.* , 1–13.
- [29] Loncan, L., Almeida, L.B., Bioucas-Dias, J.M., Briottet, X., Chanussot, J., Dobigeon, N., Fabre, S., Liao, W., Licciardi, G., Simoes, M., Tournet, J.Y., Veganzones, M., Vivone, G., Wei, Q., Yokoya, N., 2015. Hyperspectral pansharpening: a review. *IEEE Geosci. Remote Sens. Mag.* 3, 27–46.
- [30] Maas, A.L., Hannun, A.Y., Ng, A.Y., et al., 2013. Rectifier nonlinearities improve neural network acoustic models, in: *Proc. Int. Conf. Machine Learning (ICML)*, Citeseer. p. 3.
- [31] Manonmani, R., Suganya, G., 2010. Remote sensing and GIS application in change detection study in urban zone using multi temporal satellite. *Int. J. Geomatics and Geosci.* 1, 60–65.
- [32] Nascimento, J., Bioucas-Dias, J.M., 2007. Hyperspectral signal subspace estimation, in: *Proc. IEEE Int. Conf. Geosci. Remote Sens. (IGARSS)*, pp. 3225–3228.
- [33] Nascimento, J., Dias, J., 2005. Vertex component analysis: a fast algorithm to unmix hyperspectral data. *IEEE Trans. Geosci. Remote Sens.* 43, 898–910.
- [34] Nielsen, A.A., 2007. The regularized iteratively reweighted MAD method for change detection in multi-and hyperspectral data. *IEEE Trans. Image Process.* 16, 463–478.
- [35] Nielsen, A.A., Conradsen, K., Simpson, J.J., 1998. Multivariate alteration detection (MAD) and MAF postprocessing in multispectral, bitemporal image data: New approaches to change detection studies. *Remote Sens. Environment* 64, 1–19.
- [36] Niu, X., Gong, M., Zhan, T., Yang, Y., 2018. A conditional adversarial network for change detection in heterogeneous images. *IEEE Geosci. Remote Sens. Lett.* 16, 45–49.
- [37] Pepe, M.S., 2000. Receiver operating characteristic methodology. *J. Amer. Stat. Assoc.* 95, 308–311.
- [38] Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: *Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer. pp. 234–241.
- [39] Senthilkumaran, N., Vaithegi, S., 2016. Image segmentation by using thresholding techniques for medical images. *Computer Science & Engineering: An Int. Journal* 6, 1–13.
- [40] Shafique, A., Cao, G., Khan, Z., Asad, M., Aslam, M., 2022. Deep learning-based change detection in remote sensing images: A review. *Remote Sens.* 14, 871.
- [41] Sun, Y., Lei, L., Guan, D., Kuang, G., 2021a. Iterative robust graph for unsupervised change detection of heterogeneous remote sensing images. *IEEE Trans. Image Process.* 30, 6277–6291.
- [42] Sun, Y., Lei, L., Li, X., Tan, X., Kuang, G., 2020. Patch similarity graph matrix-based unsupervised remote sensing change detection with homogeneous and heterogeneous sensors. *IEEE Trans. Geosci. Remote Sens.* 59, 4841–4861.
- [43] Sun, Y., Lei, L., Li, X., Tan, X., Kuang, G., 2021b. Structure consistency-based graph for unsupervised change detection with homogeneous and heterogeneous remote sensing images. *IEEE Trans. Geosci. Remote Sens.* 60, 1–21.
- [44] Sun, Y., Xu, H., Ma, Y., Wu, M., Mei, X., Huang, J., Ma, J., 2023. Dual spatial-spectral pyramid network with transformer for hyperspectral image fusion. *IEEE Trans. Geosci. Remote Sens.* 61, 1–16.
- [45] Tewkesbury, A.P., Comber, A.J., Tate, N.J., Lamb, A., Fisher, P.F., 2015. A critical synthesis of remotely sensed optical image change detection techniques. *Remote Sens. Environment* 160, 1–14.
- [46] Vivone, G., Alparone, L., Chanussot, J., Dalla Mura, M., Garzelli, A., Licciardi, G.A., Restaino, R., Wald, L., 2014. A critical comparison among pansharpening algorithms. *IEEE Trans. Geosci. Remote Sens.* 53, 2565–2586.
- [47] Wald, L., Ranchin, T., Mangolini, M., 1997. Fusion of satellite images of different spatial resolutions: Assessing the quality of resulting images. *Photogrammetric Engineering Remote Sens.* 63, 691–699.
- [48] Wang, F., Xu, Y.J., 2010. Comparison of remote sensing change detection techniques for assessing hurricane damage to forests. *Environmental monitoring and assessment* 162, 311–326.
- [49] Wang, W., Fu, X., Zeng, W., Sun, L., Zhan, R., Huang, Y., Ding, X., 2021. Enhanced deep blind hyperspectral image fusion. *IEEE Trans. Neural Netw. Learn. Syst.* .
- [50] Wang, W., Zeng, W., Huang, Y., Ding, X., Paisley, J., 2019. Deep blind hyperspectral image fusion, in: *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pp. 4150–4159.
- [51] Wei, Q., Dobigeon, N., Tournet, J.Y., 2015a. Bayesian fusion of multi-band images. *IEEE J. Sel. Topics Signal Process.* 9, 1117–1127.
- [52] Wei, Q., Dobigeon, N., Tournet, J.Y., 2015b. Fast fusion of multi-band images based on solving a Sylvester equation. *IEEE Trans. Image Process.* 24, 4109–4121.
- [53] Xu, Q., Shi, Y., Zhu, X.X., 2023. UCDFormer: Unsupervised change detection using real-time transformers, in: *Proc. Joint Urban Remote Sensing Event (JURSE)*, IEEE. pp. 1–4.
- [54] Xu, Y., Wu, Z., Chanussot, J., Wei, Z., 2020. Hyperspectral images super-resolution via learning high-order coupled tensor ring representation. *IEEE Trans. Neural Netw. Learn. Syst.* 31, 4747–4760.
- [55] Zhang, H., Lin, M., Yang, G., Zhang, L., 2021a. Escnet: An end-to-end superpixel-enhanced change detection network for very-high-resolution remote sensing images. *IEEE Trans. Neural Netw. Learn. Syst.* .
- [56] Zhang, L., Nie, J., Wei, W., Li, Y., Zhang, Y., 2020. Deep blind hyperspectral image super-resolution. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 2388–2400.
- [57] Zhang, X., Huang, W., Wang, Q., Li, X., 2021b. Ssr-net: Spatial-spectral reconstruction network for hyperspectral and multispectral image fusion. *IEEE Trans. Geosci. Remote Sens.* 59, 5953–5965.
- [58] Zhao, N., Wei, Q., Basarab, A., Dobigeon, N., Kouamé, D., Tournet, J.Y., 2016. Fast single image super-resolution using a new analytical solution for $\ell_2 - \ell_2$ problems. *IEEE Trans. Image Process.* 25, 3683–3697.
- [59] Zhao, W., Wang, Z., Gong, M., Liu, J., 2017. Discriminative feature learning for unsupervised change detection in heterogeneous images based on a coupled neural network. *IEEE Trans. Geosci. Remote Sens.* 55, 7066–7080.