

Spatial Context Awareness for Unsupervised Change Detection in Optical Satellite Images

Lukas Kondmann¹, Aysim Toker, Sudipan Saha², *Member, IEEE*, Bernhard Schölkopf³, Laura Leal-Taixé⁴,
and Xiao Xiang Zhu⁵, *Fellow, IEEE*

Abstract—Detecting changes on the ground in multitemporal Earth observation data is one of the key problems in remote sensing. In this article, we introduce Sibling Regression for Optical Change detection (SiROC), an unsupervised method for change detection (CD) in optical satellite images with medium and high resolutions. SiROC is a spatial context-based method that models a pixel as a linear combination of its distant neighbors. It uses this model to analyze differences in the pixel and its spatial context-based predictions in subsequent time periods for CD. We combine this spatial context-based CD with ensembling over mutually exclusive neighborhoods and transitioning from pixel to object-level changes with morphological operations. SiROC achieves competitive performance for CD with medium-resolution Sentinel-2 and high-resolution PlanetScope imagery on four datasets. Besides accurate predictions without the need for training, SiROC also provides a well-calibrated uncertainty of its predictions.

Index Terms—Change detection (CD), multitemporal, optical images, unsupervised, urban analysis.

I. INTRODUCTION

CHANGE detection (CD) is at the heart of many impactful applications of remote sensing. Studying differences in land cover and land use over time with remote sensing imagery can shed light on urbanization trends [1], [2], ecosystem

Manuscript received July 14, 2021; revised October 14, 2021; accepted November 6, 2021. Date of publication November 25, 2021; date of current version February 21, 2022. This work was supported in part by the Helmholtz Association through the joint research school “Munich School for Data Science - MUDS,” in part by the Framework of Helmholtz AI under Grant ZT-I-PF-5-01—Local Unit “Munich Unit @Aeronautics, Space and Transport (MASTR),” in part by the Helmholtz Excellent Professorship “Data Science in Earth Observation-Big Data Fusion for Urban Research” under Grant W2-W3-100, in part by the European Research Council (ERC) through the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement ERC-2016-StG-714087 (*So2Sat*), and in part by the German Federal Ministry of Education and Research (BMBF) in the framework of the International Future AI Lab “AI4EO—Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond” under Grant 01DD20001. (*Corresponding author: Xiao Xiang Zhu.*)

Lukas Kondmann and Xiao Xiang Zhu are with the Remote Sensing Technology Institute, German Aerospace Center (DLR), 82234 Weßling, Germany, and also with Data Science in Earth Observation, Technical University of Munich, 80333 Munich, Germany (e-mail: lukas.kondmann@dlr.de; xiaoxiang.zhu@dlr.de).

Aysim Toker and Laura Leal-Taixé are with the Dynamic Vision and Learning Group, Technical University of Munich, 80333 Munich, Germany (e-mail: leal.taixe@tum.de; aysim.toker@tum.de).

Sudipan Saha is with Data Science in Earth Observation, Technical University of Munich (TUM), 80333 Munich, Germany (e-mail: sudipan.saha@tum.de).

Bernhard Schölkopf is with the Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany, and also with the ETH Zürich, 8092 Zürich, Switzerland.

Digital Object Identifier 10.1109/TGRS.2021.3130842

dynamics [3], surface water and sea ice trends [4], [5], and damages through natural disasters [6]–[8]. Because of rising spatial and temporal resolutions of Earth observation imagery, the possibilities of multitemporal analysis have increased significantly [9]. Combined with the open data policy of the Copernicus program, it is, for example, possible to acquire a Sentinel-2 image with 10-m resolution per pixel of any region of interest on any continent every five days [10] free of charge. Commercial providers of satellite imagery can even offer almost daily coverage with high-resolution imagery for large parts of the planet [11]. These trends emphasize the increasing opportunities in monitoring Earth from space and the relevance of CD as a field within remote sensing. Obtaining labeled data for CD, however, is costly in terms of time and effort, especially at scale. Therefore, a large focus of attention in the design of CD algorithms is unsupervised methods that do not require ground truth [12].

The applicability of unsupervised CD methods in multispectral satellite images varies depending on the spatial resolution of input images. For very-high-resolution (VHR) imagery with a spatial resolution up to 0.5 m, deep-learning-based methods tend to be in general preferable because of their elaborate capacity to model spatial context [12] although most of the work in this area focuses on supervised methods [13]–[18]. Since, for VHR imagery, an object, such as a building, consists of a number of pixels, modeling spatial context is essential to provide accurate unsupervised change segmentations. Saha *et al.* [12] introduce deep change vector analysis (DCVA), a VHR CD framework that combines ideas from image differencing with feature extraction based on pre-trained neural networks. DCVA has also been combined with self-supervised pretraining of the feature extractor specifically for remote sensing images [19]. MSDRL [20] is a scale-driven unsupervised method that uses deep feature extraction to obtain a pseudoclassification of change superpixels. Superpixels with high certainty pseudolabels are then taken as input to train a support vector machine, which eventually classifies the uncertain superpixels. Such preclassification schemes where pseudolabels are obtained based on another method have also been presented in conjunction with methods for unsupervised CD in synthetic aperture radar images [4], [21], [22], such as PCANet [23], [24]. Gong *et al.* [25] introduced modeling the difference image with a generative adversarial network (GAN). While the deep learning methods above were primarily designed for high-resolution imagery, some of them can be applied to medium-resolution imagery as well. In the case of

DCVA, there also exists a variant adjusted to the spatial and spectral scales of Sentinel-2 [26].

For medium-resolution CD, nondeep learning methods based and improved on change vector analysis (CVA) can still compete. CVA takes the difference of radiometric values or features derived from it over time [12] and applies a threshold to this difference image. Examples of features that have been derived from radiometric values as input for image differencing are vegetation indices [27] or tasseled cap transformation features [28]. Otsu thresholding [29] has been shown to be effective for thresholding the difference image [30] although a variety of approaches exist [31]–[33]. Beyond binary CD, the signal in the CVA difference image can also be used to uncover the type of change [34], [35].

CVA-based methods can still be insightful especially with medium resolution because the size of objects in these images is typically assumed to be similar to the spatial resolution of a pixel. However, extensions of CVA still fall short of the deep-learning-based DCVA for unsupervised CD on the OSCD benchmark. Still, the relative performance of traditional methods based on CVA improves with medium-resolution imagery compared to higher resolutions. More recent versions of CVA that can also be applied to higher resolution imagery tend to include close spatial context of pixels to some extent. Parcel CVA (PCVA) includes surrounding information of pixels by independent hierarchical segmentation at several scales [36]. Robust CVA (RCVA) improves on potential coregistration errors in the CVA framework by replacing a point in the difference image with the difference to a neighboring pixel if the difference to this neighbor is smaller [30]. Object CVA (OCVA) computes histograms of object sizes in an image and incorporates this information into a CVA framework [37]. Image differencing methods have also been successfully combined with morphological operations that allow transitioning from the pixel to the object level [38].

Although neighboring pixels are somewhat included in the change analysis of a pixel in these extensions of CVA, the spatial extent of incorporated information is small compared to the effective window of sequential convolutional operations in neural networks. Neighborhood in this context is defined not only as the immediate neighbors to a pixel of interest but also its larger spatial context up to a distance measure. The distant neighborhood of a pixel may help to identify changes because it is also affected by local trends in the image but unaffected by the change itself. For example, consider an explosion of a building between preimage and postimage, such as in the Beirut dataset used in the following. Analyzing the distant neighborhood allows to separate the actual change (destroyed building) from local trends, such as dust and dirt, which remains on surrounding buildings stirred by the explosion. However, the use of distant neighborhood context has only found limited application in CD thus far. This is particularly surprising since applications of image differencing in other domains, such as astronomy emphasize the importance of the relation of a pixel to its neighborhood [39].

Wang *et al.* [39] present the causal pixel model (CPM) for the study of multitemporal Kepler data that is used to spot transiting exoplanets in front of distant stars observed by the

space telescope. The method is also more generally known as half-sibling regression (HSR) [40]. Their task is conceptually related to a CD problem in remote sensing since it is also centered around spotting changes in multitemporal reflection intensities, which should be unrelated to the acquisition conditions. In their case, these deviations hint toward a transient object in front of a distant star rather than a change on the ground, but the fundamental principle is similar. They solve this task by modeling pixels as a function of their distant neighbors. With this model, it is possible to obtain a prediction for pixels in subsequent time steps based on their distant neighbors and compare the prediction to the actual value of the pixel. The size of the difference between predictions of pixels and their actual values is interpreted as the strength of the change signal.

HSR is related to the application of local binary patterns [41] for CD in more traditional image recognition problems. Bilodeau *et al.* [42] design a method based on local binary similarity patterns (LBSPs) to separate the image background from changes in multitemporal images. In their method, a binary similarity measure is computed between a pixel of interest and its closest neighbors within an image. If the binary similarity pattern updates notably between images, this is considered to be a change signal. A version of LBSPs has also been applied to CD in remote sensing where multitemporal images are split into overlapping blocks, and LBSPs of these blocks are compared across time [43]. Similarly, the graph structure of image patches across time has been used for homogenous and heterogenous CD [44]. The shared principle between LBSP and HSR is the approach to compare a pixel to its neighborhood and inspect how this relationship changes over time to discover potential changes. However, HSR relies mostly on distant neighborhood information rather than a small set of close neighbors and models this relationship explicitly to obtain a prediction for subsequent time periods.

One key property of HSR is the fact that it is by design comparably robust to registration errors and varying acquisition conditions for a given sensor [39]. This is because variations in the acquisition conditions can also affect distant neighbors in an image, whereas actual changes at the pixel level should be independent of distant context. Changing acquisition conditions and registration errors, however, are two of the primary sources of false positives (FPs) in CD [45]. Since HSR deals comparably well with these issues, it may work especially well for CD in remote sensing time-series data. In this article, we, therefore, apply HSR for CD in remote sensing. When we know from astronomy that distant spatial context can improve resilience against varying acquisition conditions, this might be particularly helpful in remote sensing CD.

We modify HSR in two major ways to apply it as SiROC for CD in remote sensing. First, we design an ensemble version of HSR based on mutually exclusive neighborhoods. Second, we make use of morphological operations to transition from pixel-level changes to object-level changes. SiROC is tested for urban CD in medium-resolution images on the Onera Satellite CD Dataset (OSCD) and high-resolution images from the Beirut Harbor Explosion of 2020 [Beirut Harbor Explosion

Dataset (BHED)]. Outside the urban context, we test SiROC on the Barrax Agriculture dataset and the Lamar Alpine dataset. Our main contributions are threefold.

- 1) We introduce SiROC, a robust method for unsupervised CD in optical remote sensing that combines ideas from HSR with ensembles over mutually exclusive neighborhoods and morphological operations.
- 2) SiROC achieves competitive performance for medium- and high-resolution unsupervised CDs with optical images.
- 3) SiROC also returns a built-in, well-calibrated uncertainty score with its change segmentation. The uncertainty measure allows distinguishing the predictions of the model by confidence, which is an important feature for pseudolabeling or detecting distribution shift.

II. METHOD

A. Half-Sibling Regression Image Differencing

The foundation of our method is HSR, which was originally developed for time-series analysis of the Kepler data in astronomy [39], [40], [46], [47]. Fig. 1 displays the intuition of HSR and how it is applied to obtain signals of changing pixels across time in three steps.

First, HSR models the pixel value of a star at time t as a linear combination of the pixel values of many other stars from the distant neighborhood of the pixel [see Fig. 1(a)]. The result of this first step is a linear coefficient for every included neighborhood pixel. In the second step [see Fig. 1(b)], predictions for the pixel at time $t + 1$ are obtained with HSR based on the neighboring pixels at $t + 1$ and the respective linear coefficients from step 1. If steps 1 and 2 are executed for the whole image, there is a prediction for any pixel at $t + 1$ and its actual value. The predicted image at $t + 1$ is then subtracted from the actual image value to obtain a change signal for every pixel [see Fig. 1(c)]. Intuitively, one expects a change in the pixels where the predictions based on the pixel's relation to its neighborhood in the past divert from the actual realization of the pixel. In the following, we elaborate on the formal definition of the three steps described. We restrict our description of HSR to the case of two time periods for simplicity since this is how we apply the method to remote sensing as well.

Step 1: Let $I_{x,y,t}$ be a pixel in a single-channel, 2-D image time series (I) at time t with coordinates (x, y) . HSR models the pixel $I_{x,y,t}$ as a linear combination of a set of distant neighbors N from I . The neighborhood set has the points $I_{i,j,t}$ as elements such that $(i, j) \in N_{x,y}$

$$I_{x,y,t} = \sum_{(i,j) \in N_{x,y}} \beta_{i,j,x,y} I_{i,j,t} + \epsilon_{x,y,t} \quad (1)$$

where $\beta_{i,j,x,y}$ is the coefficient of neighbor $I_{i,j,t}$ to model the point of interest $I_{x,y,t}$. $\epsilon_{x,y,t}$ is the residual of the model. Neighbors are chosen from the distant neighborhood because they might be subject to the same noise as the pixel of interest when they are selected to close to it. Wang *et al.* [39] require that an eligible neighbor has a distance of at least 20 pixels from the pixel of interest to be considered. This ensures that

the pixel of interest and the chosen neighbors have practically no overlap in stellar illumination. The number of neighboring pixels considered is generally large, and Wang *et al.* [39] select 4000 neighboring pixels in their original proposal of HSR to model one pixel of interest for Kepler data. Given the high temporal density of observations for each pixel (every 30 minutes) in Wang *et al.* [39], this is still solvable because the number of observed time periods exceeds the number of neighboring pixels used.

However, in the bitemporal case, where only one period is used for fitting, there are many potential combinations of β , which solves (1). We derive β as the closed form solution of the least-squares problem. It is a function of the pixel of interest $I_{x,y}$, the respective neighbor $I_{i,j}$, and the quadratic sum of all neighbors $I_{i',j'}$

$$\beta_{i,j,x,y} = \frac{I_{i,j,t}}{\sum_{(i',j') \in N_{x,y}} I_{i',j',t}^2} I_{x,y,t}. \quad (2)$$

Step 2: With the coefficients obtained in step 1, $I_{x,y,t+1}$ can be predicted as

$$\hat{I}_{x,y,t+1} = \sum_{(i,j) \in N_{x,y}} \beta_{i,j,x,y} I_{i,j,t+1}. \quad (3)$$

With the expression for β from (2), (3) can be rearranged to

$$\hat{I}_{x,y,t+1} = \frac{\sum_{(i,j) \in N_{x,y}} I_{i,j,t+1} I_{i,j,t}}{\sum_{(i,j) \in N_{x,y}} I_{i,j,t}^2} I_{x,y,t} \equiv g_{t+1} I_{x,y,t} \quad (4)$$

where g_{t+1} resembles a growth rate of the sum of pixel values in the selected neighbors from t to $t + 1$. In essence, the assumption is that, if the pixel values around $I_{x,y,t}$ increase by a factor g_{t+1} and no changes occurred at this location, $I_{x,y,t+1}$ should be close to $I_{x,y,t} g_{t+1}$. We can circumvent the explicit calculation of beta and directly obtain $\hat{I}_{x,y,t+1}$ based on (4), which is computationally efficient.

Step 3: The difference between $I_{x,y,t+1}$ and $\hat{I}_{x,y,t+1}$ is taken as the change signal for pixel $I_{x,y}$ between t and $t + 1$

$$I_{x,y,t+1} = \hat{I}_{x,y,t+1} + \epsilon_{x,y,t+1}. \quad (5)$$

After obtaining $I_{x,y,t+1}$ for all $(x, y) \in I_{t+1}$, the residual is given as the difference of the image matrices

$$\epsilon_{t+1} = \hat{\mathbf{I}}_{t+1} - \mathbf{I}_{t+1}. \quad (6)$$

Note that this is slightly different from the standard application of image differencing in CVA in multitemporal remote sensing. We do not directly take the difference of the image vectors at t and $t + 1$. Instead, we predict how the image would have looked like in $t + 1$ if the local neighborhood relations persisted. Then, we use this predicted image as input for image differencing with the actual image in $t + 1$. The extension of HSR to images with several channels is straightforward as one can directly sum the absolute values of ϵ for each channel to incorporate HSR information from all channels. Let $\epsilon_{t+1,c}$ be the residual of channel c of a multispectral image with C channels. Then, the aggregated change signal can be computed as

$$\epsilon_{t+1} := \sum_{c=1}^C |\epsilon_{t+1,c}|. \quad (7)$$

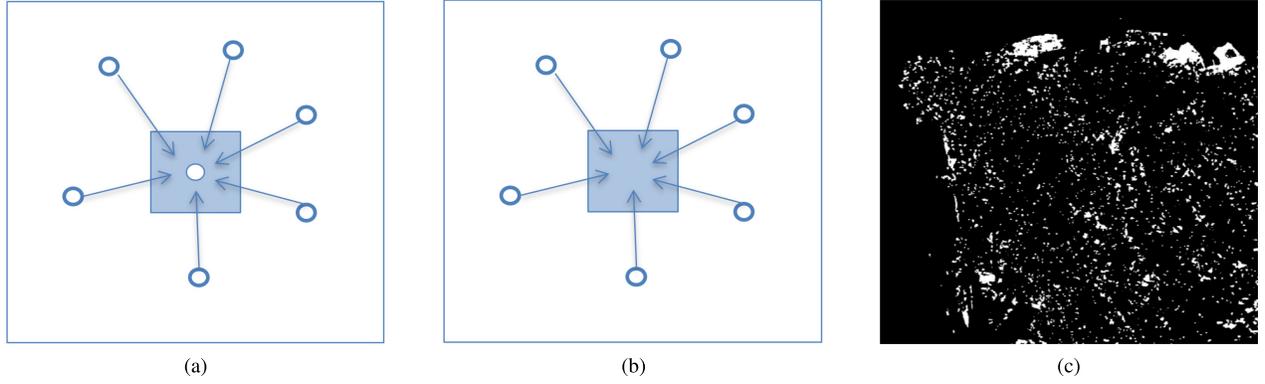


Fig. 1. HSR. (a) Set of neighbors is fit to a pixel of interest as a linear combination at time t . (b) At $t + 1$, the pixel values of the neighbors are used together with the coefficients obtained at time t to predict the pixel of interest in $t + 1$. (c) Predicted pixel values are compared with the actual pixel values at $t + 1$ to obtain a change signal.

Algorithm 1 SiROC

Input: $I_t, I_{t+1}, s, n_max, e_start$

Output: Binary Change Segmentation

```

1:  $e = e\_start, n = e\_start + s$ 
2:  $Uncertainty\_CM = zeros\_like(I_t)$ 
3: while  $n < n\_max$  do
4:   for (channel in channels) do
5:     for (pixel in  $I_t$ ) do
6:       Apply HSR( $n, e$ ) to get  $\hat{I}_{t+1}$ 
7:     end for
8:      $Channel\_Difference\_Image = \hat{I}_{t+1} - I_{t+1}$ 
9:   end for
10:   $Diff\_Image = Sum(|Channel\_Difference\_Images|)$ 
11:   $Binary\_CM = Otsu\_Thresholding(Diff\_Image)$ 
12:   $Binary\_CM\_Object = Morph\_Profile(Binary\_CM)$ 
13:   $Uncertainty\_CM = Uncertainty\_CM + Binary\_CM\_Object$ 
14:   $n = n + s$ 
15:   $e = e + s$ 
16: end while
17:  $Final\_Segmentation = Threshold(Uncertainty\_CM)$ 

```

B. HSR for Earth Observation Data (SiROC)

We improve and adapt the standard HSR image differencing model to apply it effectively for CD in remote sensing as Sibling Regression for Optical CD (SiROC). Algorithm 1 outlines SiROC in pseudocode. In summary, there are two major differences between SiROC and HSR image differencing. First, we redesign the notion of included and excluded pixels in the neighborhood selection to create an ensemble version of HSR over mutually exclusive neighborhoods. This does not only improve performance but also allows us to obtain an uncertainty along with the prediction. The rationale of splitting neighborhoods by distance is to inspect trends at different distances separately instead of pooling the trends together. For example, two trends at different spatial scales might offset each other when pooling them although both may be a signal for change. Second, we combine HSR with morphological profiles to move from pixel- to object-level

changes since changes in remote sensing typically occur at the object level.

1) *Ensembling*: The starting point for SiROC is applying HSR to I_t to obtain \hat{I}_{t+1} based on a set of neighboring pixels. We use all pixels that have a distance of at least e but at most n rows or columns from the pixel of interest. Graphically, this corresponds to all points in a square with width $2n$ and $I_{x,y,t}$ in its center, which are not in the smaller square with width $2e$ around $I_{x,y,t}$. Formally, a pixel $I_{x',y',t}$ is included in the set of neighbors for $I_{x,y,t}$ if

$$e < \max\{|x' - x|, |y' - y|\} \leq n. \quad (8)$$

With \hat{I}_{t+1} , a channel-level difference image is obtained by taking the difference $\hat{I}_{t+1} - I_{t+1}$. The absolute value of the change signal is summed across the channels. We apply Otsu-thresholding [29] to the resulting difference image that has been successfully used for thresholding difference images in CD before [30]. Furthermore, the evaluation of competing methods is also based on this thresholding approach. This allows for comparing relevant methods in a fixed setting. Nevertheless, Otsu-thresholding is a design choice here with a variety of alternatives that can also be used in conjunction with SiROC, including the T-point method [48], the Rosin method [49], or the expectation-maximization (EM) algorithm [50], [51].

The result of the thresholding step is a binary segmentation of the difference image on the pixel level. However, in remote sensing applications, changes such as the construction of roads or buildings tend to occur at the object level. This is why object-based methods often tend to be superior for these applications [52]. We rely on morphological profiles that are an established tool to bridge the gap between pixel-level change segmentations and the object level [53].

2) *Morphological Profile*: A morphological profile is the sequential application of morphological opening and closing to an image [53]. We employ morphological opening and closing at one spatial filter size p . Intuitively, morphological closing helps to fill in missed pixels in detected change objects as changed. On the other hand, morphological opening removes spurious FPs when there are no other changes around them. After obtaining an object-level change segmentation for a given neighborhood size n and exclusion

window e , we repeat the procedure and use new neighbors that are further away than the current set. Both n and e increase by the same added factor s . In the next iteration, the previous neighborhood window becomes the exclusion window, and a new binary segmentation based on more distant neighbors is obtained. This procedure is repeated until a maximum neighborhood size is reached. The number of models F is given as $F = ((n_{\max} - n_{\min}) // s) + 1$. Every model in the ensemble classifies each pixel either as change (1) or no change (0), which can be interpreted as a voting mechanism among models. Voting mechanisms across spatial scales [54] or different bands [26] are a common aggregation mechanism in CD. The number of votes per pixel ranges between 0 and F .

3) *Majority Voting*: The matrix of votes per pixel can be visualized as a heatmap of agreement between different sets of neighbors if a change occurred. This also directly transports a measure of uncertainty embedded in SiROC. If a pixel has no or the maximum number of votes, the agreement is high, and the method is confident in its prediction. If the number of votes is split, the model shows low confidence in its prediction for this point. We threshold these votes with a predefined voting share $0 \leq v \leq 1$ that is required to classify a pixel as changed. v is the sensitivity of our model toward change. The choice of v contains a tradeoff between objectives. With a higher v , the number of false negatives rises but FPs decline (and vice versa). Since all models are equally weighted in the voting process, the importance of a single neighbor is decreasing in its distance to the pixel of interest. This is because the number of neighbors used per model is increasing in n . The underlying assumption is that pixels closer to the point of interest carry more information about its potential change. This assumption is domain-specific to Earth observation and stands in contrast to the idea of HSR in astronomy where there is no weighting based on distance. The application of the voting threshold is the last step of SiROC to obtain the final change segmentation. The voting matrix is normalized by the number of models before the percentage threshold is applied.

To summarize, SiROC has the following hyperparameters.

- 1) *Maximum Neighborhood Size*: n_{\max} .
- 2) *Initial Exclusion Window*: e_{start} .
- 3) *Step Size of Ensemble*: s .
- 4) *Filter Size of Morphological Operations*: p .
- 5) *Voting Threshold*: $0 \leq v \leq 1$.

The initial size of the neighborhood window n_{start} is given as $e_{\text{start}} + s$.

III. EXPERIMENTS AND RESULTS

Section III-A describes the datasets used to assess the performance of SiROC and competing methods. The competing methods used as a benchmark and the evaluation criteria are described in more detail in Section III-B. The results on OSCD, BHED, the Agriculture Dataset, and the Alpine Dataset are presented in depth in Sections III-C–III-F, respectively.

A. Description of Datasets

1) *Onera Change Detection Dataset*: OSCD is a benchmark for bitemporal urban CD based on multispectral Sentinel-2

images [55]. It contains manual annotations of binary changes for 24 cities across the globe where 14 are used for training and 10 for testing. The labels focus on urban changes, such as newly constructed buildings, and natural changes, such as sea-level rise or differences in vegetation, are not annotated. The two images per city are selected to be cloud-free and are generally taken about one to three years apart. While there are 13 bands available in Sentinel-2 images, we restrict our focus to the RGB channels here. Although SiROC is able to handle channels outside of the visible spectrum as well, our experiments show that the inclusion of the NIR band does not add value in the urban applications considered here. This may be different in vegetation monitoring where NIR bands tend to be more insightful. Spatial bands beyond RGB and NIR do not have a spatial resolution of 10 m and are, therefore, excluded as well.

2) *Beirut Harbor Explosion Dataset*: On August 4, 2020, a devastating explosion of large amounts of ammonium nitrate occurred in the port of Beirut in Lebanon. It led to over 200 deaths and left more than 300 000 people homeless because of heavy damages to buildings in the city.¹ We collect a pair of cloud-free PlanetScope images with 3 m per pixel resolution on August 1 and 5 before and after the explosion. We combine these images with ground truth on destroyed buildings provided by the Center for Satellite Based Crisis Information (ZKI), German Aerospace Center.² The building destruction map is based on manual annotation of very-high resolution images and field reports on the ground. Note that the annotations contain building destruction rather than building damage. Therefore, partial damages to buildings that withstood the explosion are not included. With this dataset, we aim to test the applicability of SiROC not only in medium but also in higher resolution images in problems where fast and accurate annotations are essential.

3) *Agriculture Dataset*: To test SiROC also outside the urban domain, we include two other test datasets from Saha *et al.* [12] as reference points. The first one, the Agricultural dataset, is a scene with bitemporal Sentinel-2 images from July 2015 over Barrax, Spain, with 600×600 pixels in size. Between the two images is a time period of 10 days between which agricultural field activity changed notably. The reference map was manually annotated by Saha *et al.* [12].

4) *Alpine Dataset*: The second dataset consists of pre and post Sentinel-2 images of a fire in an alpine region close to Trento, Italy, in spring 2019. A variety of other seasonal vegetation trends, such as ice and snow, complicate this dataset. The scene has a size of 350×350 pixels with ground truth annotated manually by Saha *et al.* [12].

B. Competing Methods and Criteria

We compare our results to a variety of state-of-the-art unsupervised methods for optical CD in remote sensing. Since SiROC needs no training and does not rely on pretrained neural networks, its primary group of comparison consists of other image differencing-based methods. This makes SiROC

¹https://en.wikipedia.org/wiki/2020_Beirut_explosion

²<https://activations.zki.dlr.de/en/activations/items/ACT148.html>

fast compared to deep learning methods with comparable speed to traditional methods. We include several frameworks that improve on classical CVA. RCVA [30] incorporates close neighborhood information to make CVA more robust against misregistration. PCVA [36] uses CVA of multilevel parcels to improve on CVA. DCVA is based on deep feature extraction with a deep neural network pretrained on imagenet [12]. While DCVA was originally developed for high-resolution images, it is a resolution-agnostic framework relying on deep feature extraction from RGB channels. We also include a version of this method that we call DCVAMR specifically adjusted for medium-resolution, multispectral Sentinel-2 imagery [26] for the OSCD dataset. For BHED, we include DCVA, RCVA, and PCVA as baselines as DCVAMR is not capable of handling Planetscope input channels. The most recent advancement in unsupervised CD for high-resolution imagery is Saha *et al.* [19] who employ self-supervised pretraining on remote sensing images in combination with a DCVA framework. We call this refined version of DCVA “SSDCVA” and include it as the primary baseline besides general DCVA for BHED.

In line with previous evaluations on OSCD [26], [55], we analyze the performance of SiROC against the state-of-the-art binary change segmentation based on specificity and sensitivity. Specificity is defined as the number of true positives (TPs) over the sum of TPs and FPs: $\text{specificity} = \text{TN}/(\text{TN} + \text{FP})$. Sensitivity is the number of TP over the sum of TP and false negatives: $\text{sensitivity} = \text{TP}/(\text{TP} + \text{FN})$. This criterion is also known as recall. A method that is sensitive toward changes has a high sensitivity but a low specificity (and vice versa). A superior method should balance these objectives and evaluate better in both criteria. To further elaborate on the balance of change and no change class, we also report $\text{precision} = \text{TP}/(\text{TP} + \text{FP})$ and $\text{F1-score} = (2 * \text{Precision} * \text{Recall})/(\text{Precision} + \text{Recall})$.

C. Results on OSCD

1) *Parameters:* We tune the parameters of SiROC on the OSCD training set resulting in the following parameter specifications.

- 1) *Maximum Neighborhood Size:* $n_{\text{max}} = 200$.
- 2) *Initial Exclusion Window:* $e_{\text{start}} = 0$.
- 3) *Step Size of Ensemble:* $s = 8$.
- 4) *Filter Size of Morphological Operations:* $p = 5$.

The maximum neighborhood size is 200 with stepsize 8. Contrary to the original idea of HSR in astronomy, we do not find it to be optimal to exclude direct neighbors of the pixel of interest from the analysis resulting in an initial exclusion window of zero. While neighboring pixels may be subject to the same kind of object-level change on the ground, they still can contribute important information if their weight is moderate. We find the best results with morphological opening and closing with a filter size of 5. We do not tune the voting threshold because this parameter does not influence the change signal directly but rather how the method balances FPs and false negatives.

TABLE I
QUANTITATIVE RESULTS’ OSCD TEST SET

	Specificity	Sensitivity	Precision	F1
SiROC	88.31%	70.71%	24.80%	36.72%
DCVAMR	78.38%	64.63%	14.01%	23.03%
DCVA	76.96%	69.02%	14.03%	23.33%
PCVA	75.61%	47.00%	9.50%	15.81%
RCVA	76.96%	64.08%	13.16%	21.84%
Ablation Scores				
No MP	80.64%	69.88%	16.44%	26.62%
HSR	79.45%	70.24%	15.70%	25.66%

2) *Quantitative Results:* Table I reports specificity and sensitivity scores of SiROC and competing methods on the OSCD test set. Scores are averaged on the city level. SiROC with a voting threshold of $v = 1/2$ achieves a specificity of 88.31% with a sensitivity of 70.71% and 24.80% precision and 36.72% F1-score. This is a high score in all four categories by a significant margin. The difference to DCVAMR is about 6–13 percentage points (p.p.) depending on the category. DCVA achieves a sensitivity that is slightly below but close to SiROC but lacks behind in specificity, sensitivity, and F1-score by more than 10 p.p. Compared to the best results of methods without deep-learning-based feature extraction, SiROC gains about 12 p.p. in specificity, 7 p.p. in sensitivity, 12 p.p. in precision, and 15 p.p. in F1 on RCVA.

To understand the origin of the performance difference to the previous state of the art in more detail, we provide two ablation scores of SiROC. First, we remove the morphological operations in SiROC. While morphological profiles help to transition to an object-level change mask, SiROC still exceeds previous unsupervised performance without them. No MP performance improves by 2 p.p. in specificity and 5 p.p. in sensitivity versus DCVAMR and by 4 p.p. in specificity and 1 p.p. in sensitivity versus DCVA. The resulting F1-score is 3–4 p.p. higher than deep-learning-based methods and 5–10 p.p. higher than traditional methods here. To evaluate the effectiveness of ensembling, we also provide a score for a vanilla HSR with the same neighborhood size and no exclusion window. The Vanilla HSR performs slightly better but in the range of DCVA and DCVAMR with a specificity of 79.45% and a sensitivity of 70.24%. The F1-score is about 1 p.p. lower without ensemble voting.

Therefore, the majority voting mechanism is an effective tool to extract a more granular signal from the general HSR predictions. Furthermore, the use of wide spatial context pioneered in astronomy is advantageous for CD in remote sensing as well. In summary of Table I, SiROC sets a new state-of-the-art unsupervised CD in medium-resolution images on OSCD. Even without morphological filters, the method still notably outperforms previous scores, which points to a strong signal for change information in the original HSR method and the effectiveness of the majority voting mechanism. Combined with ensembling over different neighborhoods and morphological profiles, this exceeds previous quantitative results on the OSCD dataset.

3) *Qualitative Results*: The edge of SiROC compared to other unsupervised methods in the quality of change annotations for medium-resolution imagery is also visible when inspecting the predictions for specific scenes. Fig. 2 displays exemplary change masks for Las Vegas. For SiROC, a threshold of $v = 1/2$ was used to obtain the images with a specificity of 95.28%, a sensitivity of 78.75%, and a precision of 58.14% for this scene. Fig. 2(a) visualizes the confidence of SiROC in the change propensity of a pixel as a heatmap from dark purple (0% votes) to yellow (100% votes). When comparing this to the ground truth on the bottom right, one can see that change confidence is strongly associated with the occurrence of a change. Fig. 2(b) shows the binary change map after applying the threshold to the uncertainty map. Not only does SiROC pick up on the changed areas in the image but it also fits the shapes of changing buildings fairly well. The visual similarities between Fig. 2(b) and (h) are striking, especially compared to the other segmentations of competing methods. Also, before applying the morphological operations, SiROC identifies the areas of interest in the image well although the predicted mask is naturally slightly more spurious. The morphological operations help to remove these spurious changes, but the change signal in the predictions is in line with the ground truth [see Fig. 2(c)]. DCVAMR is generally able to discover the changing regions of an image but struggles to identify the shapes of changing objects and rather fits round blobs [see Fig. 2(d)]. DCVA tends to discover large changes and overestimate their size, whereas smaller changes go undetected [see Fig. 2(e)]. This might be related to the fact that DCVA was originally designed for high-resolution optical imagery in which building changes are larger in terms of pixel size. This is in line with the fact that DCVAMR, which is explicitly adjusted for Sentinel 2, tends to fit the size of changes better even though it also struggles with change shapes. PCVA and RCVA seem to extract building footprints rather than building changes here, which leads to overcrowding of the segmentation mask.

A similar picture emerges when inspecting results for Dubai in Fig. 3, which is a slightly more complex scene since the shapes of changes differ widely. SiROC detects changing regions again well but seems to struggle with the shape of changes in the upper part of the image. The newly constructed road is identified well. Consequently, the quantitative scores on this scene are slightly lower compared to the Las Vegas Scene with 86.87 % specificity, 76.61% sensitivity, and 39.14% precision. The struggles of the competing methods are similar to the Las Vegas Scene: DCVAMR fits round shapes to any kind of change [see Fig. 3(d)], DCVA overestimates the size of large changes [see Fig. 3(e)], and PCVA extracts a spurious change map that rather looks, such as building footprints [see Fig. 3(f)]. Therefore, the quality inspection of visual results confirms that SiROC obtains superior results on OSCD.

4) *Uncertainty Estimation*: To properly analyze if the confidence of SiROC also corresponds to well-calibrated uncertainties, we test this with calibration curves. For this, we split pixels into subsets based on the SiROC confidence and analyze the respective performance for a level of confidence. If the performance of SiROC is in principle increasing with the

TABLE II
SENSITIVITY TO HYPERPARAMETERS (OSCD TRAINING SET)

	Specificity	Sensitivity	Range	# Evaluations
Selected	89.48%	67.90%		
N_max	87.96%	70.53%	[30,250]	20
e_start	89.37%	68.59%	[0,150]	20
s	88.89%	68.06%	[1,5]	5
p	91.81%	59.06%	[2,5]	4
Joint	90.34%	60.12%	All above	75

Average Specificity and Sensitivity on the OSCD training set when varying SiROC hyperparameters. The average scores are compared to the selected optimal selection of parameters and their training performance in the first line. For single parameters (rows 2-5), we vary only the mentioned parameter on the grid given in the column range and leave the others at the selected optimum. The column Evaluations gives the number of runs that were executed to obtain the average scores. Finally, in the last row "Joint" we vary all parameters on the given intervals in the rows above simultaneously.

TABLE III
THRESHOLDING CHOICE (OSCD TRAINING SET)

	Specificity	Sensitivity
Otsu	89.48%	67.90%
EM	87.96%	74.22%
Triangle	88.98%	71.41%

Results on the OSCD training set with different binary thresholding techniques. We use Otsu thresholding to ensure comparability to previous results but alternatives such as expectation-maximization based thresholding give similar results.

confidence, the uncertainty levels, in fact, correspond to the certainty of the prediction that the model has. Fig. 4 plots these confidence–performance curves for four cities in the OSCD test set. For all four cities, we see that model precision is nondecreasing in the confidence of the SiROC. Most of the time, the prediction increases notably in the confidence, which means that SiROC not only performs well for this task but also returns well-calibrated uncertainties as part of its prediction.

5) *Sensitivity to Hyperparameters*: To allow effective use of SiROC in practice, we offer a sensitivity analysis of the hyperparameter choice on OSCD along with recommendations for this choice in other applications. This sensitivity analysis is executed on the training set to avoid multiple evaluations on the test set. The results are shown in Table II. While the performance of the method naturally varies with the choice of hyperparameters, SiROC looks fairly robust against its hyperparameter choices. The first row gives the training set performance based on the selected parameters described in this section as a comparison point. Varying only the maximum neighborhood N_max , the number of rows excluded e_start and the stepsize s at the selected parameter specification influences the training performance marginally, at most. For all three parameters, the average specificity decreases, while average sensitivity increases slightly. These three parameters essentially navigate how to group and prioritize neighborhoods. Excluding close context (e_start), including more distant context (N_max), and aggregating neighborhoods into larger groups (s), therefore, do not seem to matter notably in practice to achieve good performance.

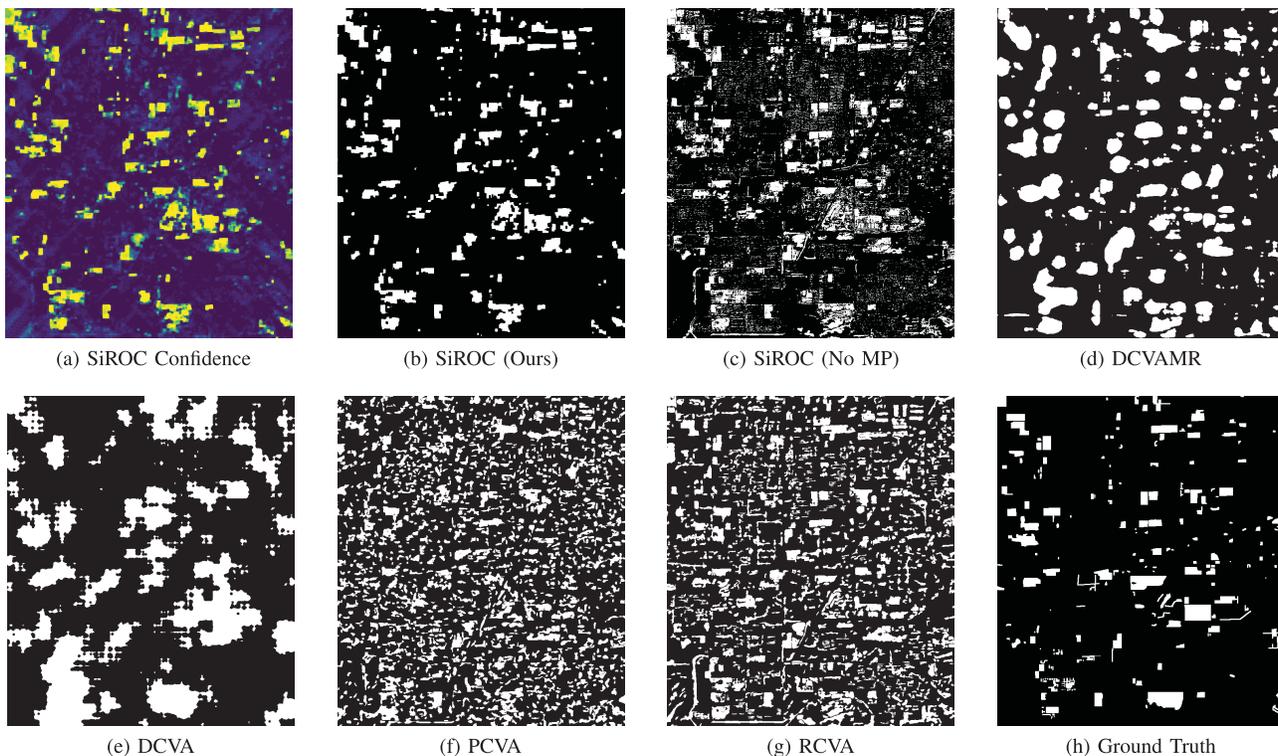


Fig. 2. Qualitative Comparison OSCD—Las Vegas. This figure visualizes the number of change votes per pixel in SiROC (a) and the corresponding binary predictions after (b) and before morphological operations (c). Competing models are visible in (d)–(g) and the ground truth in (h). SiROC predicts change regions and shapes of the ground truth well while competing methods struggle either with identifying the shapes visible in (d) and (e) or the areas of change in (f) and (g) for the Las Vegas Pair.

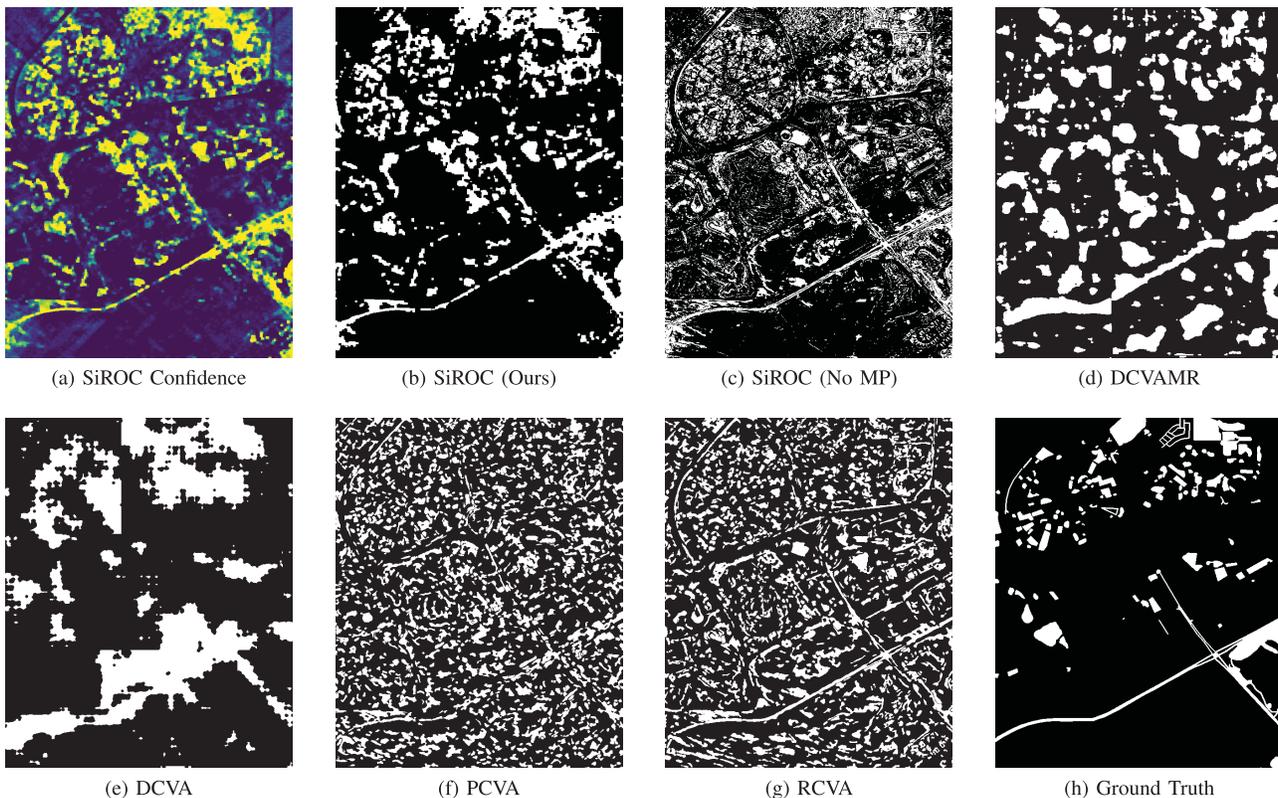


Fig. 3. Qualitative comparison OSCD—Dubai. The structure is identical to Fig. 2, but predictions and ground truth are presented for Dubai. Also, for this scene, SiROC predicts changing areas and their shape comparably well even. In contrast, competing methods miss the shapes of changing areas, such as the street in the lower part of the image or struggle to detect relevant regions.

The performance is slightly more sensitive toward the size of the morphological profile (p) where average specificity increases marginally and sensitivity drops by 9 p.p. if this is

varied leaving other parameters untouched. Similarly, when varying all four parameters simultaneously in 75 random draws, performance drops with a difference of about 8 p.p.

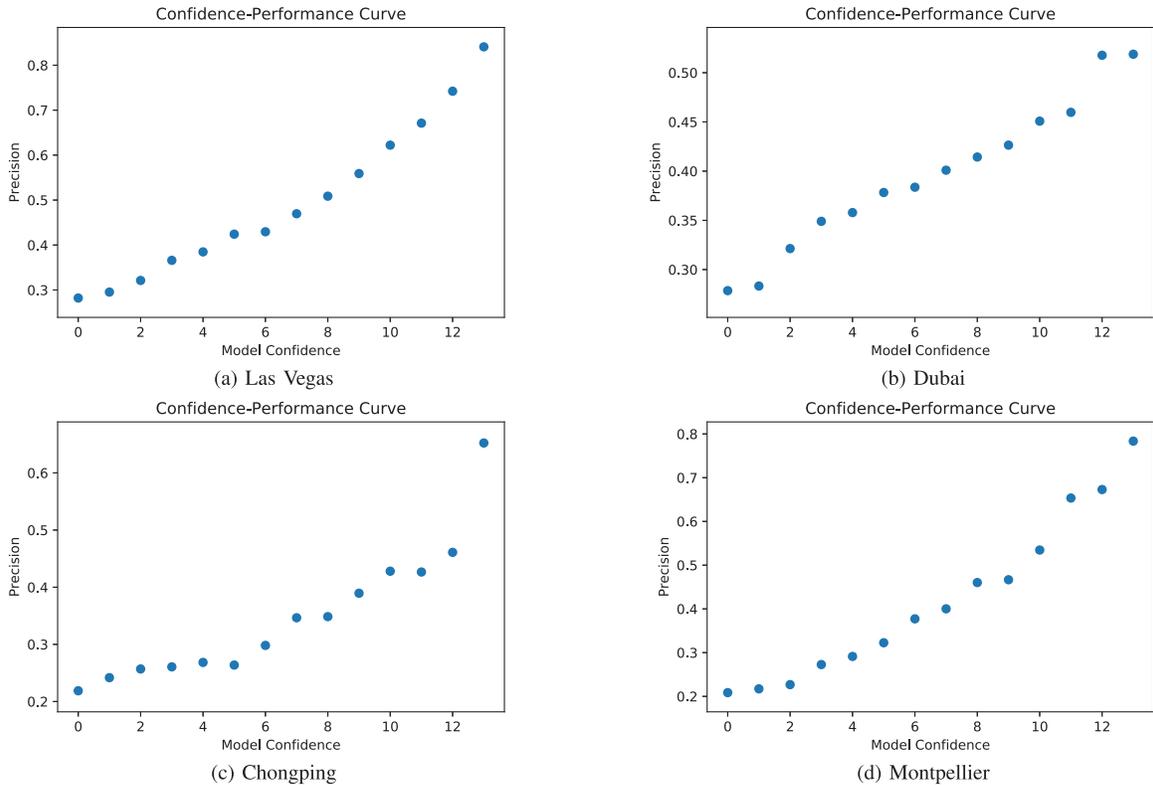


Fig. 4. Confidence–performance plots on four cities of the OSCD dataset. On the x -axis, points in the image are sorted into buckets by SiROC confidence. For each of these buckets, the performance is estimated separately. Since performance is generally nondecreasing in confidence, the uncertainty measure is considered well calibrated.

in sensitivity with similar specificity. In terms of magnitude, this performance drop is only a fraction of the difference between SiROC and its closest competitors on the OSCD test set. This implies that SiROC would likely outperform competing methods on this dataset for a variety of hyperparameter choices. To understand the sensitivity of the results to different thresholding techniques, we also benchmark SiROC based on the selected hyperparameters with EM-based thresholding [50], [51] and triangle thresholding following the OpenCV implementation.³ Results are presented in Table III. While the results are similar, the different techniques balance the tradeoff between FPs and false negatives in a slightly different fashion. This may be relevant to consider for applications of SiROC in practice where this balance plays an important role.

For potential applications of SiROC in the future, we suggest using the obtained parameter combination initially. This provides a starting point for further analysis in different contexts. Since the performance seems to be comparably susceptible to the size of the morphological profile, this parameter may deserve special attention during tuning. In the following, results on the remaining three datasets are obtained with this parameter combination, which was the result of tuning on OSCD. Even though this may not necessarily give the best possible performance, we aim to validate that SiROC achieves convincing results in other applications without fine-tuning on single scenes.

³https://docs.opencv.org/4.5.3/d7/d1b/group__imgproc__misc.html

TABLE IV
QUANTITATIVE RESULTS BEIRUT EXPLOSION

	Specificity	Sensitivity	Precision	F1
SiROC	92.01%	83.38%	19.89%	32.12%
DCVA	91.87%	79.85%	11.37%	19.93%
SSDCVA	88.25%	81.08%	8.80%	15.95%
PCVA	88.61%	58.56%	6.74%	12.10%
RCVA	86.56%	66.71%	6.52%	11.89%
Ablation Scores				
SiROC (p=10)	92.34%	91.89%	22.20%	35.76%
no MP	88.02%	79.67%	13.65%	23.30%
HSR	86.65%	71.63%	11.31%	19.54%

D. Results on BHED

1) *Quantitative Results:* Table IV displays specificity, sensitivity, precision, and F1-scores on the scene. Generally, scores on BHED are higher than on OSCD since the changes are centered around the same area and have similar shapes. SiROC with default parameters achieves a specificity of 92.01% and a sensitivity of 83.38%. DCVA achieves a similar specificity with 91.87% but falls short in terms of sensitivity by about 4 p.p. with a score of 79.85%. SSDCVA places slightly below DCVA with a specificity of 88.25% and a sensitivity of 81.08%. SiROC beats SSDCVA by about 3 p.p. in sensitivity and about 2 p.p. in sensitivity. PCVA and RCVA clearly fall behind SiROC and also DCVA-based methods. F1-score and precision results confirm the previous impressions with a gap of 12–20 p.p. in F1 and 9–13 p.p. in precision, respectively.

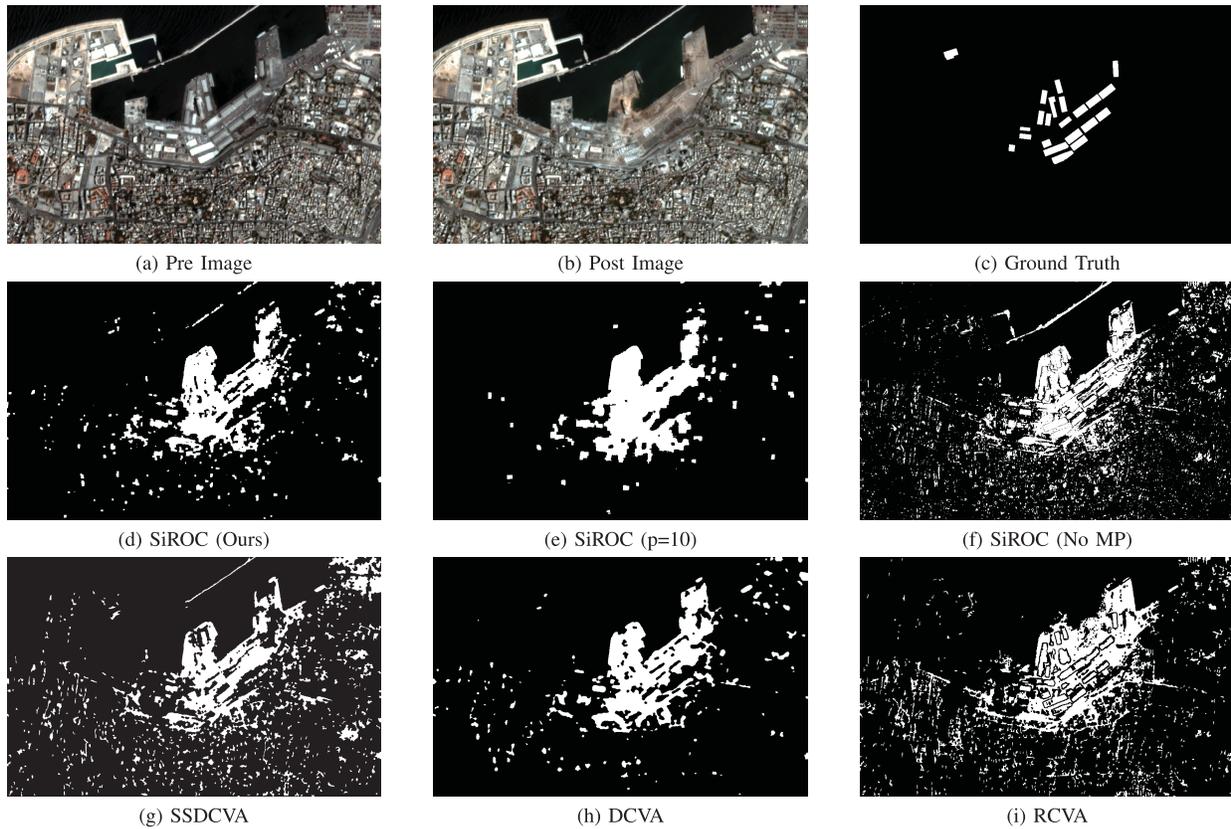


Fig. 5. Qualitative comparison Beirut explosion. This figure shows the PlanetScope image pairs (a) and (b), change ground truth (c), and model predictions (d)–(i) for the Beirut explosion scene. SiROC with main parameters (d) identifies the area of destroyed buildings around the epicenter correctly with few FPs although the shapes are lacking some granularity. Increasing the size of morphological operations improves accuracy but tends to fit one large blob with missing building shapes (e). Excluding morphological operations increases FPs although the main changes in the center are still identified well (f). Competing methods struggle not only with the shape of change but also detect a number of FPs far away from the explosion (g)–(i).

When we adjust the scale of morphological operations to 10, SiROC performs significantly better, which suggests that there may be notable tuning potential for higher resolution inputs. Still, the baseline parameters perform well on this scene. Therefore, SiROC demonstrates its usefulness beyond medium-resolution images and can also be used in conjunction with high-resolution images for CD.

The other ablation scores again point toward the most important steps within SiROC to achieve this performance. Without morphological profiles, the scores of SiROC drop about 4–9 p.p. in all four categories. Still, it achieves slightly superior precision and F1-scores but falls short to DCVA with a difference of about 3 p.p. in specificity and similar sensitivity. This is a notable difference to medium-resolution imagery on OSCD where the exclusion of morphological filters decreased the performance of SiROC, but it was still superior to DCVA-based methods. This is not necessarily surprising since deep-learning-based methods tend to relatively improve their CD performance compared to traditional methods with increasing spatial resolution. Without majority voting over different neighborhoods, the Vanilla HSR version performs better but in the range of RCVA and PCVA. Again, it is the combination of HSR, ensembling over different neighborhoods and transitioning to the object level with morphological operations that all contribute significantly to the overall performance of SiROC.

2) *Qualitative Results*: Fig. 5 shows visual comparisons of the discussed methods on BHED. The first row of images presents the preexplosion image [see Fig. 5(a)], the post image [see Fig. 5(b)], and the ground truth [see Fig. 5(c)]. The heart of the explosion in the port can be found in the middle of the image with almost the entirety of buildings completely destroyed around it. Fig. 5(d) presents the binary SiROC segmentation with baseline parameters obtained on OSCD. While SiROC is missing some granularity in its segmentation of destroyed building footprints, the changing areas are well identified with few FPs outside of the port. For a larger morphological filter size (p), the main area is identified more densely with better quantitative results, but the shapes of buildings vanish [see Fig. 5(e)]. Without morphological operations, the core change is still well-segmented although the number of FPs in the outer regions of the image increases [see Fig. 5(f)]. SSDCVA shows similar tendencies to summarize the port area as one large change with a number of spurious FPs [see Fig. 5(g)]. DCVA shows fewer salt and pepper noise than SSDCVA here and generally segments the exploded buildings similar to SiROC, however, with a slightly more perforated shape [see Fig. 5(h)]. The segmentation by RCVA is not really competitive here since the maps are spurious and changes are not well identified [see Fig. 5(i)]. Results for PCVA are similar to RCVA and, hence, omitted.

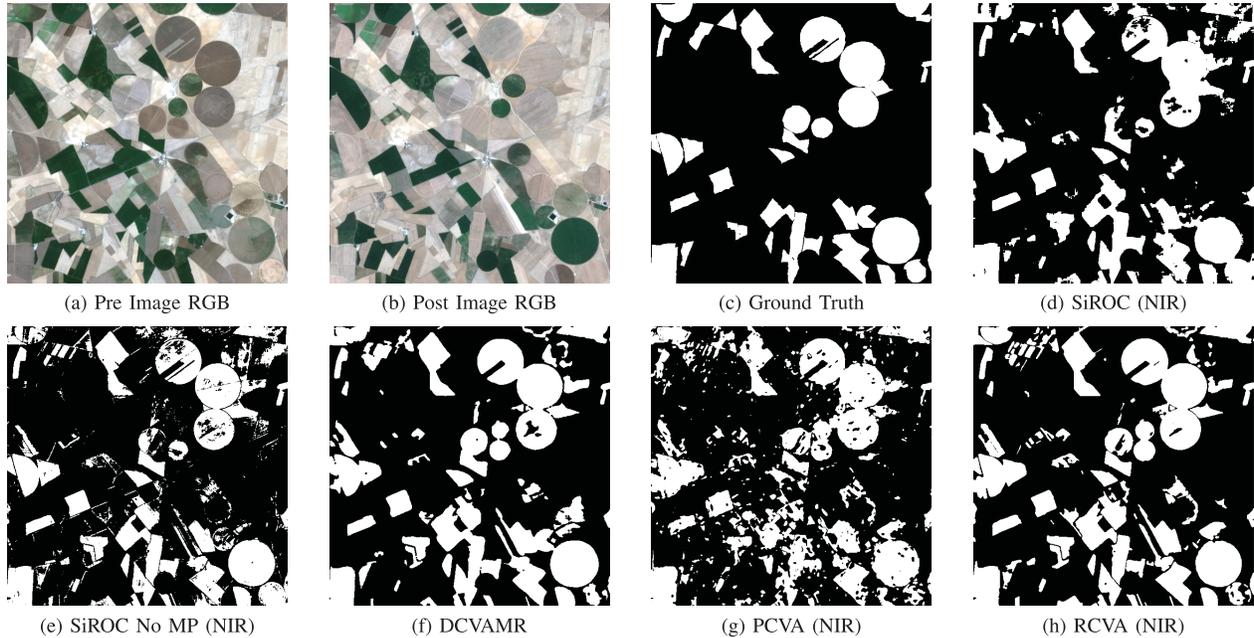


Fig. 6. Qualitative results agriculture dataset. Pre- (a) and post -GB (b) image with changing agricultural fields. The ground truth (c) shows high similarity to SiROC (d) but also to DCVA (f) and RCVA (h) while PCVA (g) has some FP areas and spurious change predictions. Change segmentations without morphological profiles (e) for SiROC still works well which is line with the quantitative results of Table V.

E. Results on Agriculture Dataset

1) *Quantitative Results:* Table V displays the results of SiROC and competing methods on the agriculture scene. SiROC is applied to the dataset with the parameters obtained on Onera without further adjustment. Hence, the results that we provide are a validation exercise in the different context of nonvisible parts of the spectrum without parameter fine-tuning.

To be consistent with previous evaluations on this dataset [12], we compare SiROC with PCVA and RCVA based on vegetation (VEG) and near-infrared (NIR) channels of Sentinel-2 as inputs. The score for DCVAMR is based on the full Sentinel-2 input images as the method was deliberately designed to incorporate all channels.

While SiROC achieves the top score in terms of specificity and precision with 90.81% and 74.23%, respectively, it falls short of DCVAMR on sensitivity (88.70% versus 94.26%) and F1-score (80.85% versus 81.47%). DCVAMR seems to lean slightly more toward the change class, whereas SiROC rather classifies a pixel as no change in unclear cases. SiROC is superior to PCVA and comparable to RCVA in performance for both VEG and NIR channels as inputs.

The ablation scores underline that morphological profiles still help although the effects are smaller than in urban applications with an average difference in about 1–2 p.p. in all four criteria. Furthermore, excluding the majority voting mechanism does not hurt performance but actually improves it slightly here. The vanilla HSR performs slightly worse but in the range of RCVA and better than PCVA on its own. Smaller benefits of including majority voting and morphological profiles could be linked to the fact that parameters for these operations were tuned in an urban RGB context. Although already quite effective, the accuracy of SiROC could likely be further improved with parameter fine-tuning.

TABLE V
QUANTITATIVE RESULTS AGRICULTURE DATASET

	Specificity	Sensitivity	Precision	F1
SiROC (VEG)	90.69%	86.38%	73.53%	79.44%
SiROC (NIR)	90.81%	88.70%	74.28%	80.85%
DCVAMR	88.88%	94.26%	71.73%	81.47%
PCVA (VEG)	88.83%	83.18%	69.04%	75.45%
PCVA (NIR)	86.60%	84.56%	65.38%	73.74%
RCVA (VEG)	88.91%	91.95%	71.28%	80.31%
RCVA (NIR)	87.39%	92.36%	68.67%	78.77%
Ablation Scores				
No MP (VEG)	89.87%	84.66%	71.44%	77.49%
No MP (NIR)	89.70%	87.15%	71.70%	78.67%
HSR (VEG)	89.96%	87.71%	72.35%	79.29%
HSR (NIR)	89.29%	90.64%	71.70%	80.06%

2) *Qualitative Results:* Fig. 6 presents pre and post RGB images [see Fig. 6(a) and (b)], the ground truth [see Fig. 6(c)], and change predictions [see Fig. 6(d)–(h)]. The visual impression of change predictions confirms the quantitative results. Predictions are fairly accurate on this scene, which suggests a comparably easy task relative to the more complex OSCD scenes. SiROC segments changing regions well and struggles with the varying field shapes only in rare instances. Similarly, the results of DCVAMR and RCVA are also fairly accurate with a slightly higher tendency to predict the change class. In comparison, the mask by PCVA produces some FP regions. Overall, SiROC shows similar performance to highly effective methods also in the agriculture domain.

F. Results on Alpine Dataset

1) *Quantitative Results:* Results for the Alpine dataset can be found in Table VI. Even though SiROC reaches the highest specificity, it does not quite pass the overall performance of

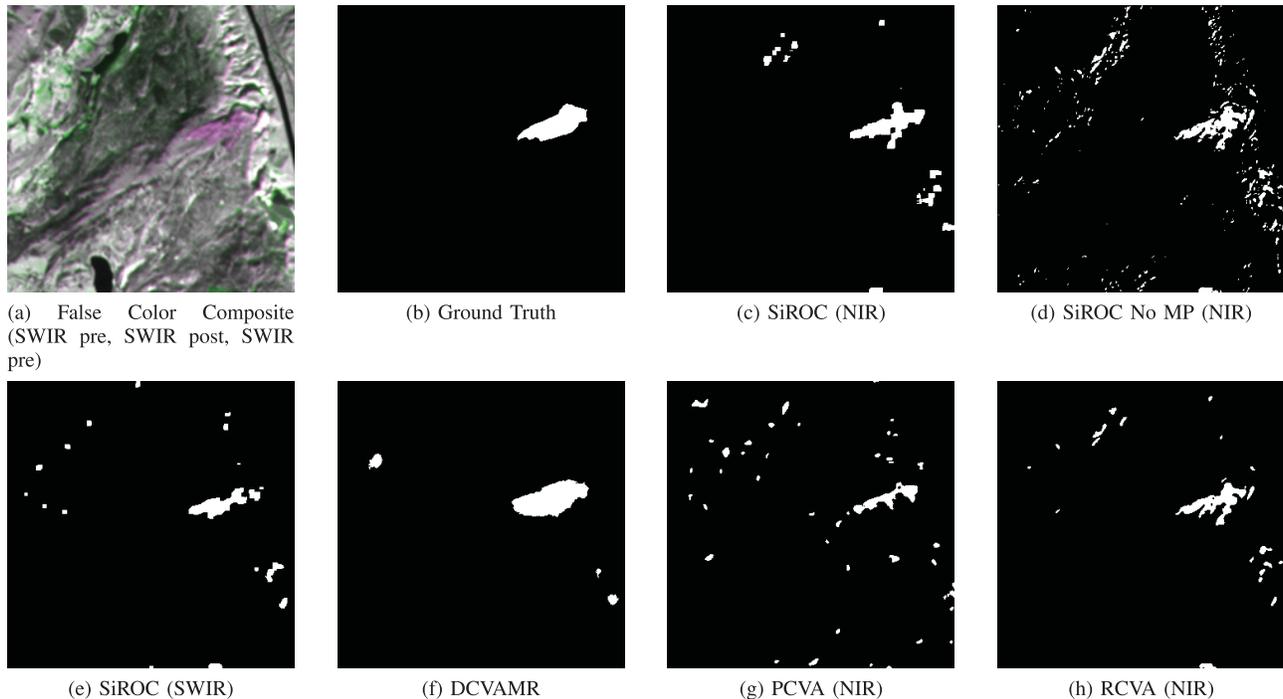


Fig. 7. Qualitative results Alpine dataset. The area of the fire can be seen in purple in the false color composite in (a) with the reference map in (b). SiROC (NIR) (c) and SiROC (SWIR) (e) both identify the changing area well although the shape is better approximated with NIR inputs. Without morphological profiles SiROC picks up more FPs (d). DCVAMR shows the most convincing results here (f). RCVA is roughly comparable (h) to SiROC while PCVA falls behind slightly (g).

DCVAMR from [12] on this dataset. Nevertheless, SiROC ranks highly also in sensitivity, precision, and F1-score, particularly based on NIR inputs with total scores of 98.92%, 75.71%, 52.28%, and 61.85%. RCVA with NIR inputs is comparable in performance, but SiROC is the only method that makes effective use of SWIR inputs compared to PCVA and RCVA.

The ablation scores underline the effectiveness of morphological transformations with about a 20 pp. drop in F1-score compared to SiROC for both NIR and SWIR. Removing the ensembling leads to a notable drop in F1-scores, particularly with NIR inputs.

2) *Qualitative Results*: Fig. 7 plots prediction masks for selected models for the Alpine dataset. In Fig. 7(a), the false color composite shows the annotated area of change affected [see Fig. 7(b)] by a fire in purple on the right. SiROC identifies this well although it is tempted to also classify a small number of FPs as change. While it is hard to control for seasonality in a bitemporal setting, SiROC (NIR) [see Fig. 7(c)] still excludes most other vegetation updates that are not the result of actual change here. The morphological profiles help on this scene to exclude spurious predictions [see Fig. 7(d)]. Compared to SiROC (SWIR) [see Fig. 7(e)], SiROC (NIR) segments the changing area slightly better although the shape is identified more clearly by DCVAMR [see Fig. 7(f)]. PCVA (NIR) [see Fig. 7(g)] seems to struggle slightly more with the shape of the burned area, whereas the results of RCVA (NIR) [see Fig. 7(h)] look similar to the results of SiROC (NIR), which is in line with the quantitative scores of Table VI.

TABLE VI
QUANTITATIVE RESULTS ALPINE DATASET

	Specificity	Sensitivity	Precision	F1
SiROC (NIR)	98.92%	75.71%	52.28%	61.85%
SiROC (SWIR)	99.28%	59.51%	56.10%	57.76%
DCVAMR	99.06%	94.99%	61.23%	74.46%
PCVA (NIR)	98.95%	46.99%	41.04%	43.82%
PCVA (SWIR)	95.48%	35.80%	10.98%	16.80%
RCVA (NIR)	99.22%	63.99%	56.20%	59.84%
RCVA (SWIR)	86.56%	66.71%	6.52%	11.89%
Ablation Scores				
No MP (NIR)	97.81%	65.58%	31.74%	42.78%
No MP (SWIR)	97.39%	55.51%	24.87%	34.36%
HSR (NIR)	95.32%	82.10%	21.45%	34.01%
HSR (SWIR)	96.32%	66.12%	21.87%	32.87%

IV. DISCUSSION

SiROC is an effective method for CD in medium- and high-resolution optical imageries, which achieves competitive performance on four datasets. In the following, we elaborate on the intuition of SiROC's performance. When contrasting SiROC to image differencing methods, SiROC can be interpreted as an improvement over standard image differencing techniques because it does not assume the same changes in the acquisition conditions across time for the whole image. Rather, it allows for local changes in acquisition conditions. In standard CVA or RCVA, for example, an implicit assumption is that changes in the acquisition conditions across time affect each pixel similarly. SiROC releases this restriction and, instead, allows for local trends in regions of the image. If a

pixel deviates from the local trend around it, it is likely to undergo a change in SiROC. In RCVA or CVA, one would compare this pixel against trends in the whole image and not against its surrounding only. This might be unrealistic in complex scenes where pixels values highly depend on local trends in the surroundings. This is, for example, the case when a new building casts a shadow on a previously illuminated pixel. Similarly, a cloudy pixel in $t + 1$ that was unobstructed in t might not necessarily be changing and is rather influenced by the local trend of a cloud rather than general image trends if large parts of the image are not obstructed by clouds. Hence, SiROC allows for a more granular analysis of deviations from trends in an image time series because, compared to previous methods, it makes full use of multitemporal information in close and distant neighbors. Although we compare our results to deep-learning-based methods, our intention is rather to augment these models than replace them, especially with high-resolution images. SiROC provides an efficient and accurate way to obtain change labels that could also be infused into deep learning models. One application of SiROC could be in self-supervised learning where pseudolabels are often obtained based on traditional image differencing techniques, such as CVA [56]. SiROC is not only superior in performance compared to image differencing. It also comes with a built-in, well-calibrated uncertainty of predictions. This could be especially beneficial in self-supervised settings since it automatically allows discriminating pseudolabels by confidence. For example, one could train only based on pseudolabels with high certainty and discard uncertain data points. Similarly, in some unsupervised methods, such as MSDRL for VHR imagery, an initial pseudoclassification is separated by confidence where high confidence examples are used for training a classifier that, subsequently, obtains predictions for leftover uncertain pixels [20]. In these methods, SiROC could also be used to obtain initial predictions and uncertainties to potentially improve not only the initial classification but maybe also the uncertainty categorization. The combination of deep-learning-based methods and SiROC may hence open up new potential for CD methods. While we restrict our focus to CD with optical images here, the framework of SiROC may be extended for applications on other multitemporal CD problems in remote sensing as well.

V. CONCLUSION

We present SiROC, an efficient and accurate unsupervised method for CD in medium- and high-resolution optical images. SiROC is inspired by HSR that is used for exoplanet search in astronomy. It models a pixel of interest in t as a linear combination of its neighbors and applies this model to $t + 1$ to obtain a prediction for the pixel based on its neighbors. The difference of the prediction for $t + 1$ and the actual pixel value in $t + 1$ is interpreted as the change signal. If the prediction is far from the actual value, trends in the neighboring pixels divert from the difference in the pixel of interest over time, which is seen as an indicator for change on the ground.

We refine and extend HSR in two major ways to apply it to optical satellite images as SiROC. First, we iterate over

several, mutually exclusive neighborhoods and apply HSR with all of these neighborhoods as input to obtain a distribution of change predictions. We combine these predictions with majority voting, which improves performance significantly and also returns a heatmap of votes per pixel, which can be interpreted as a well-calibrated uncertainty. Second, we use morphological opening and closing at one spatial filter scale to transition from pixel- to object-level predictions.

The results of SiROC are validated on four datasets. For urban CD with medium-resolution images, we verify the effectiveness of our method on OSCD, which contains binary change annotations for 24 cities across the globe. SiROC sets a new state-of-the-art unsupervised CD on OSCD, which surpasses previous methods by 10 p.p. in terms of specificity, 2 p.p. in sensitivity, 11 p.p. in precision, and 13 p.p. in F1-score. We further validate the performance of SiROC on high-resolution images with a dataset on the Beirut Harbor Explosion (BHED). Also, in this dataset, SiROC surpasses the performance of competing methods and underlines its abilities to segment urban change accurately at several scales. Furthermore, we provide two validation exercises on nonurban data with Sentinel-2 inputs. SiROC segments the effects of a fire in the Italian Alps accurately and in the range of competing methods. On the Agriculture dataset, SiROC falls short of DCVAMR in overall scores but still identifies the changing crop activity correctly.

While SiROC compares well against current deep-learning-based unsupervised methods in CD, SiROC should rather be seen as a complement than a substitute to these methods. Since it provides an accurate way to predict change signals with a built-in, well-calibrated uncertainty, it may be especially useful in conjunction with deep-learning-based methods to generate pseudolabels. Although we apply SiROC primarily to changes with multispectral data, the model may be applicable to other CD problems as well, which we plan to explore in future research.

ACKNOWLEDGMENT

The authors are grateful to the Center for Satellite Based Crisis Information (ZKI), German Aerospace Center, for providing the ground truth of the Beirut Explosion scene.

REFERENCES

- [1] D. Lu, E. Moran, and S. Hetrick, "Detection of impervious surface change with multitemporal Landsat images in an urban-rural frontier," *ISPRS J. Photogram. Remote Sens.*, vol. 66, no. 3, pp. 298–306, May 2011.
- [2] S. Ji, Y. Shen, M. Lu, and Y. Zhang, "Building instance change detection from large-scale aerial images using convolutional neural networks and simulated samples," *Remote Sens.*, vol. 11, p. 1343, Jun. 2019.
- [3] G. Chen and G. J. Hay, "An airborne lidar sampling strategy to model forest canopy height from quickbird imagery and GEOBIA," *Remote Sens. Environ.*, vol. 115, no. 6, pp. 1532–1542, Jun. 2011.
- [4] Y. Gao, F. Gao, J. Dong, and S. Wang, "Transferred deep learning for sea ice change detection from synthetic-aperture radar images," *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 10, pp. 1655–1659, Oct. 2019.
- [5] K. Rokni, A. Ahmad, K. Solaimani, and S. Hazini, "A new approach for surface water change detection: Integration of pixel level image fusion and image classification techniques," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 34, pp. 226–234, Feb. 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0303243414001780>
- [6] R. Gupta *et al.*, "Creating xBD: A dataset for assessing building damage from satellite imagery," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2019, pp. 10–17.

- [7] L. Moya *et al.*, "Detecting urban changes using phase correlation and ℓ_1 -based sparse model for early disaster response: A case study of the 2018 Sulawesi Indonesia earthquake-tsunami," *Remote Sens. Environ.*, vol. 242, Jun. 2020, Art. no. 111743.
- [8] M. Zanetti *et al.*, "A system for burned area detection on multi-spectral imagery," *IEEE Trans. Geosci. Remote Sens.*, early access, Sep. 17, 2021, doi: 10.1109/TGRS.2021.3110280.
- [9] X. X. Zhu *et al.*, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [10] M. Drusch *et al.*, "Sentinel-2: ESA's optical high-resolution mission for GMES operational services," *Remote Sens. Environ.*, vol. 120, pp. 25–36, May 2012.
- [11] C. Kwan *et al.*, "Assessment of spatiotemporal fusion algorithms for planet and worldview images," *Sensors*, vol. 18, no. 4, p. 1051, Mar. 2018.
- [12] S. Saha, F. Bovolo, and L. Bruzzone, "Unsupervised deep change vector analysis for multiple-change detection in VHR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 6, pp. 3677–3693, Jun. 2019.
- [13] Y. Zhan, K. Fu, M. Yan, X. Sun, H. Wang, and X. Qiu, "Change detection based on deep Siamese convolutional network for optical aerial images," *IEEE Geosci. Remote Sens. Lett.*, vol. 14, no. 10, pp. 1845–1849, Oct. 2017.
- [14] H. Lyu, H. Lu, and L. Mou, "Learning a transferable change rule from a recurrent neural network for land cover change detection," *Remote Sens.*, vol. 8, no. 6, p. 506, 2016.
- [15] L. Mou, L. Bruzzone, and X. X. Zhu, "Learning spectral-spatial-temporal features via a recurrent convolutional neural network for change detection in multispectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 2, pp. 924–935, Feb. 2019.
- [16] P. Zhang, M. Gong, L. Su, J. Liu, and Z. Li, "Change detection based on deep feature representation and mapping transformation for multi-spatial-resolution remote sensing images," *Photogram. Remote Sens.*, vol. 116, pp. 24–41, Sep. 2016.
- [17] S. Saha, L. Mou, X. X. Zhu, F. Bovolo, and L. Bruzzone, "Semisupervised change detection using graph convolutional network," *IEEE Geosci. Remote Sens. Lett.*, vol. 18, no. 4, pp. 607–611, Apr. 2021.
- [18] M. Gong, Y. Yang, T. Zhan, X. Niu, and S. Li, "A generative discriminatory classified network for change detection in multispectral imagery," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 1, pp. 321–333, Jan. 2019.
- [19] S. Saha, L. Mou, C. Qiu, X. X. Zhu, F. Bovolo, and L. Bruzzone, "Unsupervised deep joint segmentation of multitemporal high-resolution images," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 12, pp. 8780–8792, Dec. 2020.
- [20] T. Zhan, M. Gong, X. Jiang, and M. Zhang, "Unsupervised scale-driven change detection with deep spatial-spectral features for VHR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 8, pp. 5653–5665, Aug. 2020.
- [21] M. Gong, H. Yang, and P. Zhang, "Feature learning and change feature classification based on deep learning for ternary change detection in SAR images," *J. Photogramm. Remote Sens.*, vol. 129, pp. 212–225, Jul. 2017.
- [22] F. Gao, X. Wang, Y. Gao, J. Dong, and S. Wang, "Sea ice change detection in SAR images based on convolutional-wavelet neural networks," *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 8, pp. 1240–1244, Aug. 2019.
- [23] F. Gao, J. Dong, B. Li, and Q. Xu, "Automatic change detection in synthetic aperture radar images based on PCANet," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 12, pp. 1792–1796, Dec. 2016.
- [24] M. Li, M. Li, P. Zhang, Y. Wu, W. Song, and L. An, "SAR image change detection using PCANet guided by saliency detection," *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 3, pp. 402–406, Mar. 2018.
- [25] M. Gong, X. Niu, P. Zhang, and Z. Li, "Generative adversarial networks for change detection in multispectral imagery," *IEEE Geosci. Remote Sens. Lett.*, vol. 14, no. 12, pp. 2310–2314, Nov. 2017.
- [26] S. Saha, Y. T. Solano-Correa, F. Bovolo, and L. Bruzzone, "Unsupervised deep transfer learning-based change detection for HR multispectral images," *IEEE Geosci. Remote Sens. Lett.*, vol. 18, no. 5, pp. 856–860, May 2021.
- [27] A. Singh, "Review article digital change detection techniques using remotely-sensed data," *Int. J. Remote Sens.*, vol. 10, no. 6, pp. 989–1003, 1989.
- [28] Y. T. S. Correa, F. Bovolo, and L. Bruzzone, "Change detection in very high resolution multisensor images," in *Proc. 20th Image Signal Process. Remote Sens.*, vol. 9244, 2014, Art. no. 924410.
- [29] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-9, no. 1, pp. 62–66, Jan. 1979.
- [30] F. Thonfeld, H. Feilhauer, M. Braun, and G. Menz, "Robust change vector analysis (RCVA) for multi-sensor very high resolution optical satellite data," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 50, pp. 131–140, Aug. 2016.
- [31] L. Bruzzone and D. F. Prieto, "Automatic analysis of the difference image for unsupervised change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 38, no. 3, pp. 1171–1182, May 2000.
- [32] L. Bruzzone and D. F. Prieto, "A minimum-cost thresholding technique for unsupervised change detection," *Int. J. Remote Sens.*, vol. 21, no. 18, pp. 3539–3544, 2000.
- [33] T. Celik, "Unsupervised change detection in satellite images using principal component analysis and k -means clustering," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 4, pp. 772–776, Aug. 2009.
- [34] F. Bovolo and L. Bruzzone, "A theoretical framework for unsupervised change detection based on change vector analysis in the polar domain," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 1, pp. 218–236, Jan. 2006.
- [35] F. Bovolo, S. Marchesi, and L. Bruzzone, "A framework for automatic and unsupervised detection of multiple changes in multitemporal images," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 6, pp. 2196–2212, May 2011.
- [36] F. Bovolo, "A multilevel parcel-based approach to change detection in very high resolution multitemporal images," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 1, pp. 33–37, Jan. 2008.
- [37] L. Li, X. Li, Y. Zhang, L. Wang, and G. Ying, "Change detection for high-resolution remote sensing imagery using object-oriented change vector analysis method," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2016, pp. 2873–2876.
- [38] N. Falco, G. Cavallaro, P. R. Marpu, and J. A. Benediktsson, "Unsupervised change detection analysis to multi-channel scenario based on morphological contextual analysis," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2016, pp. 3374–3377.
- [39] D. Wang, D. W. Hogg, D. Foreman-Mackey, and B. Schölkopf, "A causal, data-driven approach to modeling the Kepler data," *Publications Astronomical Soc. Pacific*, vol. 128, no. 967, 2016, Art. no. 094503.
- [40] B. Schölkopf *et al.*, "Modeling confounding by half-sibling regression," *Proc. Nat. Acad. Sci. USA*, vol. 113, no. 27, pp. 7391–7398, Jul. 2016.
- [41] M. Heikkilä, M. Pietikäinen, and C. Schmid, "Description of interest regions with local binary patterns," *Pattern Recognit.*, vol. 42, no. 3, pp. 425–436, 2009.
- [42] G.-A. Bilodeau, J.-P. Jodoin, and N. Saunier, "Change detection in feature space using local binary similarity patterns," in *Proc. Int. Conf. Comput. Robot. Vis.*, May 2013, pp. 106–112.
- [43] N. Gupta, G. V. Pillai, and S. Ari, "Change detection in optical satellite images based on local binary similarity pattern technique," *IEEE Geosci. Remote Sens. Lett.*, vol. 15, no. 3, pp. 389–393, Mar. 2018.
- [44] Y. Sun, L. Lei, X. Li, X. Tan, and G. Kuang, "Patch similarity graph matrix-based unsupervised remote sensing change detection with homogeneous and heterogeneous sensors," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 6, pp. 4841–4861, Jun. 2021.
- [45] L. Bruzzone and F. Bovolo, "A conceptual framework for change detection in very high resolution remote sensing images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Jul. 2010, pp. 2555–2558.
- [46] D. Wang, D. W. Hogg, D. Foreman-Mackey, and B. Schölkopf, "A pixel-level model for event discovery in time-domain imaging," 2017, *arXiv:1710.02428*.
- [47] T. D. Gebhard, M. J. Bonse, S. P. Quanz, and B. Schölkopf, "Physically constrained causal noise models for high-contrast imaging of exoplanets," Dec. 2020, *arXiv:2010.05591*. [Online]. Available: <https://arxiv.org/abs/2010.05591>
- [48] N. Coudray, J.-L. Buessler, and J.-P. Urban, "Robust threshold estimation for images with unimodal histograms," *Pattern Recognit. Lett.*, vol. 31, no. 9, pp. 1010–1019, Jul. 2010.
- [49] P. L. Rosin, "Unimodal thresholding," *Pattern Recognit.*, vol. 34, no. 11, pp. 2083–2096, 2001.
- [50] Y. Bazi, L. Bruzzone, and F. Melgani, "Image thresholding based on the EM algorithm and the generalized Gaussian distribution," *Pattern Recognit.*, vol. 40, no. 2, pp. 619–634, Feb. 2007.
- [51] M. Zanetti, F. Bovolo, and L. Bruzzone, "Rayleigh-Rice mixture parameter estimation via EM algorithm for change detection in multispectral images," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 5004–5016, Dec. 2015.

- [52] A. Song, Y. Kim, and Y. Han, "Uncertainty analysis for object-based change detection in very high-resolution satellite images using deep learning network," *Remote Sens.*, vol. 12, no. 15, p. 2345, Jul. 2020.
- [53] M. D. Mura, J. A. Benediktsson, F. Bovolo, and L. Bruzzone, "An unsupervised technique based on morphological filters for change detection in very high resolution images," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 433–437, Jul. 2008.
- [54] S. Liu, Q. Du, X. Tong, A. Samat, L. Bruzzone, and F. Bovolo, "Multi-scale morphological compressed change vector analysis for unsupervised multiple change detection," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 10, no. 9, pp. 4124–4137, Sep. 2017.
- [55] R. C. Daudt, B. L. Saux, A. Boulch, and Y. Gousseau, "Urban change detection for multispectral Earth observation using convolutional neural networks," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2018, pp. 2115–2118.
- [56] H. Dong, W. Ma, Y. Wu, J. Zhang, and L. Jiao, "Self-supervised representation learning for remote sensing image change detection based on temporal prediction," *Remote Sens.*, vol. 12, no. 11, p. 1868, Jun. 2020.



Lukas Kondmann received the bachelor's degree in economics from the Ludwig Maximilian University of Munich, Munich, Germany, in 2016, the Honors degree in technology management from the Center for Digital Technology and Management, Munich, in 2017, and the master's degree in social data science from the University of Oxford, Oxford, U.K., in 2019. He is currently pursuing the Ph.D. degree in engineering with the Technical University of Munich, Munich, and the German Aerospace Center, Munich.

He was a Visiting Researcher working on big data for social good with the School of Information, University of California at Berkeley (UC Berkeley), Berkeley, CA, USA, in spring 2017. His research is centered around time-series analysis of multispectral remote sensing imagery with a focus on monitoring the sustainable development goals (SDGs).



Aysim Toker received the B.Sc. degree in computer engineering and the M.Sc. degree in computer science from the Technical University of Munich, Munich, Germany, in 2016 and 2018, respectively, where she is currently pursuing the Ph.D. degree with the Dynamic Vision and Learning Group, under the supervision of Prof. Dr. Laura Leal-Taixe for a joint project with Prof. Dr. Xiaoxiang Zhu.

She has been a programmer for many years, with broad experience in many languages. Her interest in computer vision and machine learning started with her M.Sc. degree. During her master's thesis, she concentrated on video object segmentation. Currently, she is doing research in deep learning, sequence analysis, and remote sensing with a focus on the intersection of these three domains.



Sudipan Saha (Member, IEEE) received the M.Tech. degree in electrical engineering from IIT Bombay, Mumbai, India, in 2014, and the Ph.D. degree in information and communication technologies from the University of Trento, Trento, Italy, and the Fondazione Bruno Kessler, Trento, in 2020.

He was an Engineer with TSMC Ltd., Hsinchu, Taiwan, from 2015 to 2016. In 2019, he was a Guest Researcher with the Technical University of Munich (TUM), Munich, Germany, where he is currently a Post-Doctoral Researcher. His research interests are related to multitemporal remote sensing image analysis, domain adaptation, time-series analysis, image segmentation, deep learning, image processing, and pattern recognition.

Dr. Saha was a recipient of the Fondazione Bruno Kessler Best Student Award in 2020. He is a reviewer for several international journals. He has served as a Guest Editor for *Remote Sensing* (MDPI) special issue on "Advanced Artificial Intelligence for Remote Sensing: Methodology and Application."



Bernhard Schölkopf has researched at AT&T Bell Labs, Holmdel, NJ, USA, GMD FIRST, Berlin, Germany, and Microsoft Research Cambridge, Cambridge, U.K., becoming a Max Planck Director in 2001. He is currently a Professor with ETH Zürich, Zürich, Switzerland. He co-initiated the MLSS series of Machine Learning Summer Schools, the Cyber Valley Initiative, and the ELLIS grassroots initiative. He has applied his methods to a number of different fields, ranging from biomedical problems to computational photography and astronomy. His scientific interests are in machine learning and causal inference.

Dr. Schölkopf is also a fellow of the Association for Computing Machinery (ACM) and the CIFAR Program Learning in Machines and Brains and a member of the German Academy of Sciences. He (co)received the Academy Prize of the Berlin-Brandenburg Academy of Sciences and Humanities, the Royal Society Milner Award, the Leibniz Award, the Koerber European Science Prize, and the BBVA Foundation Frontiers of Knowledge Award. He is the Co-Editor-in-Chief of the *Journal of Machine Learning Research*, an early development in open access and today the field's flagship journal.



Laura Leal-Taixé received the B.Sc. and M.Sc. degrees in telecommunications engineering from the Technical University of Catalonia (UPC), Barcelona, Spain, in 2005 and 2008, respectively, and the Ph.D. degree from the Leibniz University of Hannover, Hanover, Germany, in 2013.

She was a Visiting Scholar with the University of Michigan, Ann Arbor, MI, USA. She spent two years as a Post-Doctoral Researcher at ETH Zürich, Zürich, Switzerland, and a year as a Senior Post-Doctoral Researcher at the Computer Vision Group, Technical University of Munich, Munich, Germany. She went to Boston, MA, USA, to do her master's thesis at Northeastern University, Boston, with a fellowship from the Vodafone foundation. She is currently a tenure-track Professor (W2) with the Technical University of Munich, leading the Dynamic Vision and Learning Group.

Dr. Leal-Taixé was a recipient of the Sofja Kovalevskaja Award of 1.65 million euros for her project socialMaps and the Google Faculty Award.



Xiao Xiang Zhu (Fellow, IEEE) received the M.Sc., Dr.Eng., and "Habilitation" degrees in signal processing from the Technical University of Munich (TUM), Munich, Germany, in 2008, 2011, and 2013, respectively.

She was a Guest Scientist or a Visiting Professor with the Italian National Research Council (CNR-IREA), Naples, Italy, Fudan University, Shanghai, China, The University of Tokyo, Tokyo, Japan, and the University of California at Los Angeles, Los Angeles, CA, USA, in 2009, 2014, 2015, and 2016, respectively. Since 2019, she has been a Co-Coordinator of the Munich Data Science Research School, Munich. Since 2019, she has been heading the Helmholtz Artificial Intelligence—Research Field "Aeronautics, Space, and Transport." Since May 2020, she has been the Director of the international future AI lab "AI4EO—Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond," Munich. Since October 2020, she has been the Co-Director of the Munich Data Science Institute (MDSI), TUM. She is currently a Professor with Data Science in Earth Observation (former: Signal Processing in Earth Observation), TUM, and the Head of the Department "EO Data Science," Remote Sensing Technology Institute, German Aerospace Center (DLR), Weßling, Germany. She is currently a Visiting AI Professor with ESA's Phi-Lab, Frascati, Italy. Her main research interests are remote sensing and Earth observation, signal processing, machine learning, and data science, with a special application focus on global urban mapping.

Dr. Zhu is also a member of the young academy, Junge Akademie/Junges Kolleg, at the Berlin-Brandenburg Academy of Sciences and Humanities, the German National Academy of Sciences Leopoldina, and the Bavarian Academy of Sciences and Humanities. She also serves on the scientific advisory board in several research organizations, including the German Research Center for Geosciences (GFZ) and the Potsdam Institute for Climate Impact Research (PIK). She is also an Associate Editor of IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING. She also serves as the Area Editor responsible for special issues of *IEEE Signal Processing Magazine*.