

Monitoring of Construction Activity by Change Detection on SAR Time Series Using Coherent Scatterers

Carlos Villamil Lopez  and Uwe Stilla , *Senior Member, IEEE*

Abstract—In this article, a novel synthetic aperture radar (SAR) change detection method for the monitoring of man-made objects (MMOs) is presented. Rather than looking for changes in SAR amplitude or the loss of coherence, changes are detected by the appearance and disappearance of the strong point scatterers present in MMOs and often denoted as coherent scatterers (CSs). This enables the detection of changes involving MMOs while ignoring changes to natural targets such as vegetation. These CSs are detected in each image and compared coherently across an image pair or a time series. When using a time series, the proposed method can categorize the changes according to their temporal behavior. An object-based change analysis step for identifying changes significantly larger than individual CSs is also introduced. The proposed approach is applied for the monitoring of construction activity using a time series of 49 TerraSAR-X images of the city of Munich.

Index Terms—Change detection (CD), coherent scatterers (CSs), monitoring, synthetic aperture radar (SAR), time series.

I. INTRODUCTION

SPACEBORNE synthetic aperture radar (SAR) sensors are especially well suited to monitor changes on the ground because, unlike optical sensors, they can operate in adverse weather conditions and independently of sunlight [1]. When comparing two SAR images acquired with the same sensor and imaging geometry (i.e., using a repeat-pass orbit) at different times, any significant differences between them will be due to changes in the imaged scene. SAR missions with mid-resolution sensors like Sentinel-1 provide a global coverage every few days and can be exploited to monitor changes of large or moderate size on a regional or global scale. On the other hand, missions with high-resolution capabilities, such as TerraSAR-X or COSMO-SkyMed, can regularly acquire much more detailed images of specific locations, enabling the detection of smaller changes. Change detection (CD) techniques exploiting these

high-resolution SAR images can be applied to monitor anthropogenic objects, also often called man-made objects (MMOs). Changes caused by their appearance, disappearance, or movement inside the imaged scene can be detected, as well as changes to static objects. This enables the monitoring of different types of human activity, such as the arrival and departure of airplanes at airports [2] and of ships at ports [3], the construction of new buildings [4], the movement of shipping containers and parked cars [5], or changes in the amount of oil stored in refineries [6].

CD methods with SAR images can be classified into coherent and incoherent [7], [8]. Incoherent change detection (ICD) methods detect changes by comparing the amplitude of two coregistered SAR images, while coherent change detection (CCD) methods detect changes by the loss of coherence. CCD methods can detect subtle changes [7], such as those caused by vehicles when driven over soft surfaces [9] or by objects when displaced to a distance smaller than the pixel size. Those changes would not typically be resolved with the amplitude of the SAR images. However, CCD can only be applied to data acquired with a short temporal baseline to 1) avoid decorrelation [10] and 2) reduce false alarms caused by natural targets [11]. ICD methods are less sensitive, but they can be applied to image pairs with longer temporal baselines and in some cases even with slightly different imaging geometries [12].

Traditional ICD and CCD methods cannot easily distinguish different types of changes, such as those caused by MMOs or seasonal changes like vegetation growth or snow. Seasonal changes are irrelevant for many applications, but they are often detected as they can induce significant amplitude changes and coherence loss. For example, snow can dominate the CD results even for short temporal baselines, making the resulting change maps of little use [11]. Methods affected by this are not very well suited for applications focusing only on MMOs, as seasonal changes would result in false alarms. This also applies to modern CD methods using time series [13], [14], [15], polarimetric SAR data [16], [17], or deep learning approaches trained to detect changes in SAR amplitude [18], [19], [20].

In addition, conventional CCD and ICD methods cannot unambiguously distinguish changes due to the appearance and disappearance of MMOs. For CCD methods, both types of changes cause the coherence to drop significantly. ICD methods can distinguish changes caused by an increase or decrease in the SAR amplitude. It is often assumed that a strong increase/decrease in the SAR amplitude indicates the appearance/disappearance of an

Manuscript received 27 April 2022; revised 12 July 2022; accepted 16 August 2022. Date of publication 24 August 2022; date of current version 13 September 2022. (Corresponding author: Carlos Villamil Lopez.)

Carlos Villamil Lopez is with the Microwaves and Radar Institute, German Aerospace Center, 82234 Wessling, Germany, and also with the Unit “Photogrammetry and Remote Sensing,” Department of Aerospace and Geodesy, TUM School of Engineering and Design, Technical University of Munich, 80333 Munich, Germany (e-mail: carlos.villamillopez@dlr.de).

Uwe Stilla is with the Unit “Photogrammetry and Remote Sensing,” Department of Aerospace and Geodesy, TUM School of Engineering and Design, Technical University of Munich, 80333 Munich, Germany (e-mail: stilla@tum.de).

Digital Object Identifier 10.1109/JSTARS.2022.3201480

object. However, this assumption is not always valid, as objects often cast a shadow area where the amplitude decreases. This effect has been exploited for detecting changes associated with buildings [4].

Some recent publications have presented SAR CD methods focused on the detection of changes caused by MMOs and not those associated with natural targets, such as vegetation. These methods tackle this issue in different ways, such as by implementing hand-crafted detectors for specific types of MMOs [5], by using persistent scatterer interferometry (PSI) [21] to detect new buildings in large time series [22] or by decomposing the images in a time series into background and strong scatterers components and using the extracted scatterers for CD [23], [24]. While these approaches partially solve the aforementioned issue, they also have certain limitations. The method presented in [5] relies on hand-crafted features and needs to be tuned for different objects. The application of PSI by Yang and Soergel [22] requires a very long time series and can only detect very slow or permanent changes caused by objects that remain unchanged along many images. The method proposed in [23] also requires a time series and can at most detect a single change event per pixel during the time series. Finally, the approach described in [24] detects changes by an increase or decrease in the number of strong scatterers inside a given radius and appears unable to detect changes where the scatterers change but their number remains stable.

In this article, a novel CD approach is presented for the monitoring of MMOs using pairs or series of high-resolution single-pol SAR images. Rather than looking for changes in the SAR amplitude or the loss of coherence, these changes are detected by the appearance and disappearance of the strong point scatterers present in MMOs and often denoted as coherent scatterers (CSs) [25]. The CSs are detected in each image by analyzing their phase over different frequency subbands [26], [27] and then compared coherently. An object-based analysis step is applied to extract information on an object level for changes significantly larger than individual CSs. The proposed approach is unsupervised, and it avoids many of the previously mentioned limitations of other CD methods.

- 1) It detects only the changes associated with MMOs and ignores the changes to natural targets.
- 2) It can exploit CCD even with large temporal baselines, as CSs are mostly unaffected by temporal decorrelation.
- 3) It can distinguish the changes caused by an object's appearance, disappearance, or modification.
- 4) It works with as few as two images, and it can detect up to n different CSs per pixel in a time series with n images.
- 5) It can identify and ignore irrelevant transient changes (i.e., an object temporarily affected by an external factor).
- 6) It can target specific types of changes (e.g., by their size and/or temporal behavior) and segment them.

The rest of this article is organized as follows. Section II presents the proposed CD method. The data used to evaluate the proposed method and the experiments performed to select the method's parameters are described in Section III. Section IV shows the results obtained when applying the proposed approach to a long time series for the monitoring of construction activity.

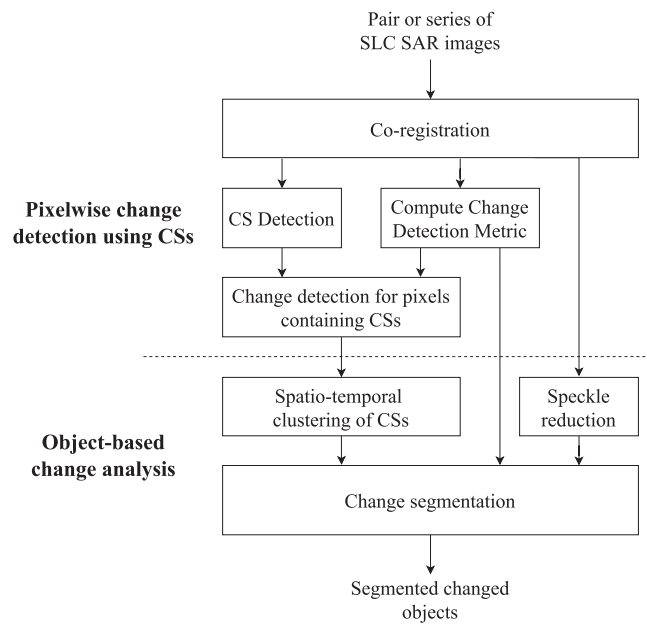


Fig. 1. Processing chain of the proposed CD approach.

Finally, a discussion is given in Section V, and future work is outlined in Section VI.

II. METHOD

In this section, we present the proposed method, named Change Detection by Coherent Scatterers. Initially, an overview of existing algorithms to detect CSs in a single-look complex (SLC) SAR image is provided in Section II-A. Then, in Section II-B, we describe how different types of changes between a pair of SAR images can be detected by coherently comparing the CSs detected in the two images. An extension of this method for longer time series is described in Section II-C. Here, we introduce a new CD metric that exploits the full coherence matrix and is able to ignore irrelevant transient changes. Because the changes to be detected are often significantly larger than individual CSs, the pixelwise CD using CSs is followed by an object-based change analysis step. This consists of a clustering step, described in Section II-D, followed by the segmentation of the changed objects, described in Section II-E. Finally, in Section II-F, we briefly describe how certain types of changes can be identified by their size and/or temporal behavior. A block diagram of the complete processing chain is shown in Fig. 1.

A. Detection of CSs in an SLC SAR Image

The strong point scatterers typically present in MMOs are often denoted as CSs [25]. CSs can be detected in an SLC high-resolution SAR image by exploiting the fact that they remain stable across multiple sublooks [25], [26], [27], [28]. Sublooks can be computed along range (i.e., different frequency subbands) or azimuth (i.e., different subapertures), as described in [25]. Because many CSs exhibit a nonconstant azimuth angular scattering pattern [28], sublooks are typically computed along range for CS detection.

Different CS detection methods have been evaluated and compared in [25]. In this work, CSs are detected by analyzing the variation of the phase with respect to sublook frequency [26]. If a pixel contains a CS, its phase varies linearly; otherwise, it varies randomly. The phase (given in radians) of a given pixel at sublook i (with $i = 1, \dots, n$) can be denoted as ϕ_i . If this phase varies almost linearly, the phase difference between each two consecutive sublooks, $\Delta\phi_i = \phi_{i+1} - \phi_i$, will be nearly constant and, therefore, have a low variance. Phase jumps of 2π around $\pm\pi$ should be compensated when computing $\Delta\phi_i$ (e.g., by applying phase unwrapping). The variance σ_ϕ^2 of the $n - 1$ samples of $\Delta\phi$ can be computed for each pixel and thresholded to detect CSs: the pixels with $\sigma_\phi^2 < T$ are considered to contain a CS. The value of T trades off the number of false positives and false negatives. For the CD method presented in this article, we recommend selecting a value of T resulting in few false positives.

The lower range resolution of the sublook images causes strong point scatterers to spread to neighboring pixels. These pixels often also exhibit a linear phase and can be detected when thresholding σ_ϕ^2 . This effect can be mitigated by analyzing the slope of the linear phase trends. This slope is proportional to the distance in range between a pixel and the corresponding CS [25]. A constraint on the slope can be analytically derived and implemented: of the pixels with $\sigma_\phi^2 < T$, only those with $\phi_n - \phi_i \leq \pi$ actually contain a CS in the full resolution image. Combining these two constraints allows detecting CSs with virtually no resolution loss.

B. CD on an Image Pair Using CSs

CCD can be used to determine whether a CS moved or changed between two SAR images acquired with the same imaging geometry at different times. If a CS experiences even a small change (e.g., a subpixel displacement), the interferometric coherence of the corresponding pixel will drop significantly. On the other hand, if this CS remains unchanged and static, the coherence will have a high value, as strong point scatterers typically have high coherence and are not significantly affected by temporal decorrelation [21], [29].

To implement such a CD method, the two SLC SAR images should first be coregistered with subpixel accuracy by using a method such as the one described in [30]. Afterward, the CSs can be detected for each image, as described in Section II-A. The pixels of the resulting two binary images get the value 0 or 1 depending on the absence or presence of a CS. The absence of a CS is assumed to be clutter. The CD metric is the coherence computed with the image pair of interest. A threshold γ_t should then be applied to the resulting coherence image in order to detect which of these CSs have changed or moved during the interval between the two image acquisitions. Finally, depending on the results of the CS detection and the coherence thresholding, different types of changes can be distinguished. The possible cases are summarized in Table I and described in the following paragraphs.

If a pixel containing a CS has a coherence below γ_t , it will imply that a change involving an MMO has very likely occurred there. Three different types of changes can then be distinguished.

TABLE I
INTERPRETATION OF THE CD RESULTS

Image 1	Image 2	Coherence	Type of change
Clutter	Clutter	–	Irrelevant
CS	Clutter	Low	Object disappeared
Clutter	CS	Low	Object appeared
CS	CS	Low	Different object or object changed/moved
CS	CS	High	Same object (no change)
CS	Clutter	High	Same object (no change)
Clutter	CS	High	Same object (no change)

If a CS was only detected in the first (i.e., earlier) image, it will imply that an object left the scene, whereas if it was detected only in the second image, it will imply that a new object appeared. If a CS was detected in both the images, then either one object changed or moved or it was replaced by a different object. On the other hand, if the coherence of a pixel containing a CS is higher than γ_t , it will be considered that this CS remained unchanged and is present in both the images, even if it was only detected in one. The reasoning behind this is that a high coherence value should only be possible if there is no significant change, whereas false negatives in the CS detection are much more likely to occur.

Finally, the pixels containing clutter in both the images are ignored, independently of the change in amplitude or coherence loss. These pixels could correspond to an environmental change, and the focus of the proposed method is to detect only those changes associated with MMOs. Besides, the proposed CCD metric would not work well for these pixels, as it assumes that unchanged pixels have high coherence values. The distributed scatterers present in natural targets are affected by temporal decorrelation. This is especially relevant at the X-band [10], which is used by most high-resolution spaceborne SAR sensors. In contrast, CCD can be robustly applied to CSs even with long temporal baselines.

When applying this method, the computed coherence image should have the same pixel spacing as the original SLC images, independently of the number of looks used to estimate the coherence and the associated resolution loss. The coherence values are only evaluated on the pixels containing a CS in at least one image. Because the images with the detected CSs have the full resolution, the window size used for the coherence estimation does not affect the spatial resolution of the results. The effect of the window size used for the coherence estimation is discussed later in Section III-B.

C. CD on a Time Series Using CSs

If a time series with more than two repeat-pass images is available, the proposed CD method can be extended. Instead of detecting changes between each of the consecutive image pairs separately, the CD task can be formulated in a slightly more general way. The goal is now to estimate the time interval $[t_{\text{start}}, t_{\text{end}}]$ when each of the detected CSs is present in the imaged scene and remains unchanged. Using CD, the exact values of t_{start} and t_{end} for a given CS cannot be determined, but these can be narrowed down to the following intervals: $t_{\text{start}} \in (t_{a-1}, t_a)$ and

$t_{\text{end}} \in (t_b, t_{b+1})$. Here, a and b are the indices of the first and last images where that specific CS was detected, and t_i denotes the acquisition time of the i th image of the series. When using a time series with n images, up to n different CSs can be detected in a given pixel, each present during a different time interval. A higher temporal resolution of the time series allows us to detect more and faster changes and better determine the change occurrence.

When performing CD with a time series, instead of simply comparing each image to the next one, all the image pairs in the series can be compared. This can be done using the coherence matrix for each pixel. This matrix contains the coherence value for images j and k (denoted as $\gamma_{j,k}$) at row j and column k . The additional coherence values can be exploited to distinguish relevant changes to an object (e.g., an object appears, leaves the scene, moves, or is modified in a lasting manner) from irrelevant transient changes. We define as transient changes those situations where an object is just temporarily affected by some external factor in one or a few outlier images and does not actually change. Possible examples are an object temporarily covered with snow or occluded due to the radar shadow casted by another object. One could argue that these situations are actual changes that should be detected, but they are typically irrelevant for many practical applications. Therefore, the ability to distinguish transient changes (e.g., a building temporarily covered with snow) from those where the object itself changes (e.g., a new building is built or an existing building is renovated or demolished) represents a clear advantage. Later in Section III-D, a real example of a transient change is shown and compared to another example where an object itself actually changes. Both are illustrated by their characteristic coherence matrices.

To determine whether a change happened to a given CS during the time interval (t_i, t_{i+1}) , we can take into account all the image pairs with one image acquired at t_i or earlier and another acquired at t_{i+1} or later. When an MMO experiences a significant change, the coherence should be low for all these image pairs. On the other hand, for transient changes, the coherence is low for the pairs containing one of the outlier images, but high for other image pairs with longer temporal baselines. This indicates some event briefly affected the object causing its coherence to drop, but the object later returned to its exact previous state, and therefore, it did not actually change.

Based on this, we can define a new CD metric f_i to determine whether a significant change happened to a given CS during the time interval (t_i, t_{i+1}) . For the case of CD with an image pair introduced in Section II-B, this metric was simply the coherence of the pair: $f_i = \gamma_{i,i+1}$. In this case, we want the new metric to be insensitive to transient changes. This can be achieved by exploiting additional image pairs with different temporal baselines, enforcing that the coherence must be low for all of them in order to detect a change. This new metric can be computed as follows:

$$f_i = \max_{\substack{j,k \\ j \leq i \\ k \geq i+1}} \gamma_{j,k}. \quad (1)$$

As this metric computes the maximum of many coherence images for each pixel, it is important to avoid noisy coherence

estimates. For this, a large window can be used during the coherence computation. Storing the coherence matrix for every pixel requires lots of memory, especially for long time series. However, this is not required: it is enough to initialize $n - 1$ empty images to store the values of f_i . After computing the coherence for a given image pair, the values of the images for f_i that include that pair can be updated accordingly. This is more memory efficient, but the coherence still needs to be computed for $(n^2 - n)/2$ image pairs. As a small subset of the coherence matrix suffices to identify transient changes, the metric f_i from (1) can be slightly modified to reduce the computational cost

$$f_i = \max_{\substack{j,k \\ j \in [i-r, i] \\ k \in [i+1, i+1+r]}} \gamma_{j,k} \quad (2)$$

where r can be used to increase the number of elements of the coherence matrix to be taken into account. A value of $r = 1$ should be enough to detect transient changes affecting only one image, whereas a value of $r = 0$ would reduce this metric to the previous one used for an image pair: $f_i = \gamma_{i,i+1}$.

To perform CD with a time series using this new metric, the CS detection should be performed for all the coregistered SLC images, as described in Section II-A. This results in a stack of binary images C_i (with i in $1, \dots, n$) with the detected CSs for each image. In addition, the CD metric should be computed using (2), resulting in a stack of images with the values of f_i for each pixel, with i in $1, \dots, n - 1$. The remaining steps can be performed exclusively for the list of pixels containing CSs instead of using the full image raster. This reduces the required memory and makes the computations faster. For this, we can compute a binary mask showing the pixels containing a CS in at least one image: $C_{\text{max}} = \max_i C_i$.

As introduced in Section II-B, some pixels might exhibit inconsistencies, with the CD metric indicating that no change occurred between two images but the CS detection indicating that only one contains a CS (i.e., $f_i \geq \gamma_t$ and $C_i \neq C_{i+1}$). Before, we argued that the most likely cause for this is false negatives in the CS detection. The proposed solution was to consider that an unchanged CS is present in both the images, effectively correcting the value of C_i or C_{i+1} . The same principle can be applied now, but the comparison should be carried over across the whole series (e.g., a CS detected in just one image could be present in several images before and after). This consistency check can be implemented with a forward pass (sequentially comparing images i and $i + 1$, with $i = 1, \dots, n - 1$), followed by a backward pass (comparing images i and $i - 1$, with $i = n, \dots, 2$). This results in corrected values for the CS detection, denoted as C'_i . These will have the same values as C_i except for the likely false negatives corrected, where $C'_i = 1$ and $C_i = 0$.

After this consistency check, the CD can finally be performed. First, the unchanged and static CSs can be identified by finding the pixels, where f_i is always above the threshold: $\min_i f_i \geq \gamma_t$. Each of these pixels contains a single CS with $t_{\text{start}} < t_1$ and $t_{\text{end}} > t_n$. For the remaining pixels, we need to determine how many different CSs were present and their corresponding time periods. For a given pixel, we can establish that a new CS appeared at a time $t_{\text{start}} \in (t_{i-1}, t_i)$, with $i > 1$, if a CS is present in image i and a change happened in between images

$i - 1$ and i (i.e., if $C'_i = 1$ and $f_{i-1} < \gamma_t$). In addition, CSs present in the first image ($C'_1 = 1$) appeared before the start of the series: $t_{\text{start}} < t_1$. In a similar way, we can establish that a CS disappeared at a time $t_{\text{end}} \in (t_i, t_{i+1})$, with $i < n$, if $C'_i = 1$ and $f_i < \gamma_t$. Also, those with $C'_n = 1$ are still present at the end of the series: $t_{\text{end}} > t_n$. After checking these conditions for all the pixels and all the images, a list of detected CSs (each with an estimated time period) is obtained for each pixel.

The proposed consistency check does not account for false positives in the CS detection. These are unlikely and usually have a negligible effect, resulting in a few wrongly detected CSs that often appear isolated, spread across the imaged scene. In contrast, the much higher number of correctly detected CSs appears in the areas where MMOs are located. However, false positives can pose a problem in certain areas that remain unchanged and are unaffected by temporal decorrelation. At these locations, the proposed consistency check would propagate the CSs wrongly detected in each image to other images of the series, resulting in a higher number of false positives. This issue becomes more significant with longer time series, but it can be mitigated by a simple postprocessing step. If the CD determines that a given CS is present in images a through b (both included), we can check the corresponding values of C_i to see in how many images it was originally detected. If this number is too low, then the corresponding CS was likely a false positive and can be discarded

$$\sum_{i=a}^b C_i < k(b - a + 1) \Rightarrow \text{discard CS, likely false positive} \quad (3)$$

with k being a value between 0 and 1, which controls the maximum fraction of false negatives to be corrected by the consistency check.

D. Spatiotemporal Clustering of CSs

A clustering algorithm can be applied to detect objects from a set of point scatterers, as MMOs appear in SAR images as clusters of densely packed CSs. The density-based spatial clustering of applications with noise (DBSCAN) algorithm [31] seems well suited for this task. Given a set of points, DBSCAN clusters together those closely packed in high-density regions. Isolated points in low-density regions are marked as outliers, which makes the algorithm robust against noise. DBSCAN can detect clusters of arbitrary shapes and its definition of a cluster aligns well with our particular problem and data. Its robustness against noise can discard most of the false positives in the CS detection. Besides, the number of clusters (unknown in our case) does not need to be specified.

The DBSCAN algorithm has two parameters: a radius ε , and p , the minimum number of points within this radius required to form a cluster. The metric used to compute the distance between points can also be specified. For our use case, we use the distance in meters between two pixels. This makes the selection of the radius ε more intuitive and compensates different pixel spacings in azimuth and range. This distance is computed by scaling the pixel coordinates of each CS by the

pixel spacing along the corresponding axes. For the range axis, typically in slant-range projection, the equivalent pixel spacing in ground-range is computed using the mean incidence angle.

The values of ε and p should simply control the density of CSs required to form a cluster. Constraints related to the minimum cluster size can be imposed later by evaluating the number of CSs and/or the area inside the convex hull of the resulting clusters. Otherwise, high values could be selected for both the parameters, making it difficult to correctly cluster objects of arbitrary shapes and increasing the probability of grouping nearby objects into the same cluster.

So far, we have explained how spatial clustering can be performed on a set of CSs. However, the temporal information obtained from the CD can also be considered for the clustering. All the CSs belonging to the same object are expected to appear and disappear at the same time. This can be enforced by splitting all the CSs into multiple subsets according to the different time intervals (i.e., the combinations of t_{start} and t_{end}). The DBSCAN algorithm can then be applied separately to each of these smaller point sets. This way, CSs closely located in an image but belonging to different objects can be separated if they have a different t_{start} or t_{end} . Besides, it also makes the clustering significantly faster, as most DBSCAN implementations run on quadratic time.

E. Segmentation of the Detected Changes

The previous clustering step grouped the CSs that are likely to belong to the same objects. However, it would be desirable to obtain a dense change map for each object. For this, we can try to segment each changed object using all the image pixels and not only CSs. The segmentation is performed separately for each cluster using all the relevant images, as t_{start} and t_{end} are known. The extent of the image patch to be processed can be obtained by getting the rectangle enclosing the cluster's convex hull. This extent can be increased by some scaling factor as a safety margin, in case the cluster is smaller than the corresponding change to be segmented.

For objects present in more than one image, the CD metric f_i can be thresholded to obtain an accurate segmentation of the changes. If an object was first seen at image a and last seen at image b , its pixels must:

- 1) change in the interval (t_{a-1}, t_a) : $f_{a-1} < \gamma_t$;
- 2) change in the interval (t_b, t_{b+1}) : $f_b < \gamma_t$;
- 3) remain unchanged in the interval $[t_a, t_b]$: $\min_{a \leq i < b} f_i \geq \gamma_t$.

These constraints result in three binary images that can be combined using a logical AND operation to obtain a change mask. This mask contains all the pixels inside the cluster and near its boundaries exhibiting the same temporal behavior. In the case that $a = 1$, it will not be possible to apply the first constraint, whereas if $b = n$, the same will happen to the second constraint. However, in both the cases, the two remaining constraints are sufficient.

For objects present in only one image (i.e., if $a = b$), it will not be possible to apply the third constraint. The other two might not suffice to segment the corresponding change, as the object

surroundings might also have low coherence due to temporal decorrelation. In this case, some additional constraints related to the SAR amplitude should be added:

- 1) amplitude change in (t_{a-1}, t_a) : $|A_a - A_{a-1}| \geq \Delta_A$;
- 2) amplitude change in (t_a, t_{a+1}) : $|A_a - A_{a+1}| \geq \Delta_A$;
- 3) amplitude should not have low values at t_a : $A_a \geq A_{\min}$.

A_i denotes the SAR amplitude in decibel (dB) scale of the image i in the series. The parameter Δ_A represents the minimum amplitude difference in dB to consider that a pixel changed. This is equivalent to the well-known log-ratio metric for ICD [32]. While there are more advanced metrics, this fixed threshold should already perform well, as it is only applied locally where changes have already been detected. The second parameter A_{\min} represents the minimum amplitude to consider that a pixel might belong to an MMO, as these typically exhibit high-amplitude values. Before applying these new constraints, the speckle noise should be reduced by applying multilooking or a more sophisticated speckle filter. Some modern methods [33], [34] achieve a very good denoising performance and preserve the full spatial resolution. These constraints only work well if the used SAR sensor is well calibrated, which is typically the case for modern spaceborne SAR sensors.

Finally, the change mask obtained for each cluster is refined by applying mathematical morphology operations. First, a closing operation with a radius of a few pixels is used to fill small holes in the mask without significantly changing its shape or size. Then, the connected components that are too small or mostly outside of the cluster's bounds are discarded. The remaining connected components should provide a good segmentation of the changed object (or objects) in the cluster.

F. Change Analysis

To detect changes corresponding to specific events, the obtained results can be analyzed to identify objects with certain temporal behaviors and of certain sizes. This is especially interesting for urban areas and similarly complex scenes where many changes occur between two consecutive image acquisitions. To focus on objects above or below certain sizes, a threshold can be applied to the area of the segmented changes. Changes can also be categorized according to their duration (e.g., into fast, long-term, and permanent changes). Also, further constraints regarding their time of appearance and disappearance can be imposed. In the following, we provide some examples of how this analysis can be applied.

Newly constructed or renovated buildings and infrastructure typically imply the appearance of new CSs, which later remain unchanged over a long time. Such changes can be identified by imposing the following constraints: $t_{\text{start}} > t_1$, $t_{\text{end}} > t_n$, and $t_{\text{end}} - t_{\text{start}} > \Delta T$. This requires at least three images: one acquired before the construction work is finished, and two acquired afterward and at least ΔT time apart. Here, ΔT is set to the minimum amount of time that an object must remain unchanged to be considered a new static object (e.g., a couple of months).

Moving objects (e.g., parked cars, airplanes in airports, etc.) typically imply CSs appearing and disappearing inside short periods: $t_{\text{end}} - t_{\text{start}} < \Delta T$. Again, at least three images are




Fig. 2. Image patch showing a building. (a) SAR amplitude. (b) Detected CSs. The range direction is indicated with a white arrow.

required to unambiguously identify this behavior: one acquired before the object appears, one with it present, and one after it leaves. Here, ΔT is set to the maximum amount of time that moving objects are expected to remain static (e.g., from multiple days to a couple of weeks).

For this kind of temporal analysis, it is important to note that t_{start} and t_{end} (and therefore also the duration) can only be narrowed down to some intervals, as introduced in Section II-C. In this article, the threshold on this time length is applied to the lower bound of this interval. However, this could be handled differently depending on the application.

III. DATA AND EXPERIMENTS

To evaluate the CD method presented in this article, a dataset consisting of 49 TerraSAR-X repeat-pass images of the city of Munich was used. These images have been acquired between March 28, 2016 and February 28, 2019; using the Staring Spotlight imaging mode (with a resolution of 58 cm in slant range and 23 cm in azimuth) and with an incidence angle of 37.5° and in ascending orbit. The SAR images shown in this article were rotated (so that the image's y -axis corresponds to range) and resized to achieve a square pixel spacing in slant range, as this allows us to better observe the layover. However, the processing is always performed using the original SLC images.

In the following, we use these data to illustrate different parts of the proposed method and to select suitable values for its parameters.

A. CS Detection

For the TerraSAR-X data with 300-MHz bandwidth used in this article, ten equally spaced sublooks with a 75% spectral overlap appear to be a good choice for CS detection. This results in a sublook bandwidth of 92.30 MHz and a spacing of 23.08 MHz. Using more sublooks with less bandwidth results in less CSs being detected in layover regions (e.g., building façades), as neighboring CSs interfere with each other in the sublook images due to their lower resolution. For the threshold T , values between 0.1 and 0.125 work well for the chosen number of sublooks. The value of T trades off the number of false positives and false negatives. An example of the CSs detected in a TerraSAR-X image of a building using a threshold $T = 0.125$ is shown in Fig. 2. The selected parameters result in a good




Fig. 3. Coherence of an image pair computed using two different window sizes. (a) 3×7 pixels. (b) 9×23 pixels. The image patch marked in yellow was used for estimating the coherence statistics for clutter listed in Table II.

TABLE II
COHERENCE STATISTICS FOR DIFFERENT WINDOW SIZES

Window size	Looks	Clutter		Point scatterers	
		Mean	Std. dev.	Mean	Std. dev.
3×7	21	0.313	0.145	0.935	0.070
5×11	55	0.213	0.105	0.906	0.087
7×17	119	0.155	0.080	0.876	0.093
9×23	207	0.124	0.065	0.861	0.102

detection performance. Most of the point scatterers visible in the SAR image are correctly detected, and almost no CSs are detected in the clutter areas. Also, the image with the detected CSs has virtually the same resolution as the input SLC image.

B. Coherence Calculation and Thresholding

The coherence plays an important role in the proposed CD method. The window size used for the coherence estimation and the threshold γ_t are important parameters. Here, we briefly analyze the effect of both and select suitable values for them. Fig. 3 shows the coherence for an image pair computed using two different window sizes. This image pair has a temporal baseline of 22 days. The coherence map in Fig. 3(a) was computed using a smaller window of 3×7 pixels and is clearly noisier. The coherence map in Fig. 3(b) was computed using a window of 9×23 pixels and has a lower resolution, but also significantly less noise. In both the cases, the window is bigger along azimuth, as the data have a higher resolution along this axis. The window sizes are chosen to achieve a similar resolution in slant range and azimuth. Table II shows the mean and standard deviation of the coherence for clutter and point scatterers, computed using four different window sizes. The small homogeneous patch highlighted in yellow in Fig. 3 was used for estimating the clutter statistics, whereas for the point scatterers, 100 of them were manually selected across the image. Table II shows that for CSs, even relatively small windows result in good coherence estimates with low bias and variance. However, for clutter and other areas with low coherence, the coherence estimates using small windows have a high bias and variance. Because of this, using small window sizes can lead to problems when applying a threshold in areas of low coherence (e.g., coherence can be overestimated where changes occurred). Therefore, in this article,

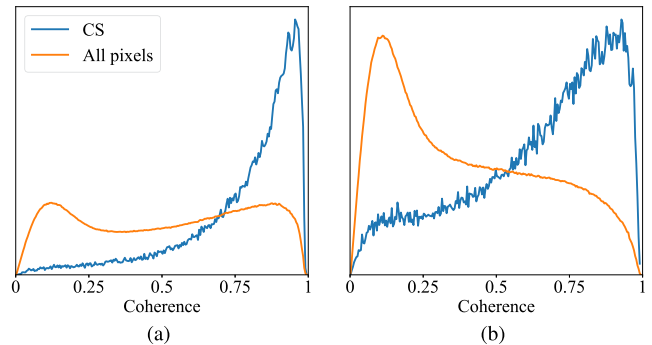


Fig. 4. Histogram of the coherence values for the CSs (blue line) and all pixels (orange line) in a large image patch in the city center. These are shown for two temporal baselines: (a) 22 days and (b) 2 years and 11 months. The histograms are normalized so that the area under the curve integrates to 1.

we use a window of 9×23 for the coherence estimation. This results in coherence maps with a resolution of approximately 5.2 m in azimuth and slant range. Nevertheless, as described in Section II-B, this does not affect the spatial resolution of the results of the pixelwise CD using CSs.

For the coherence threshold, we select a value of $\gamma_t = 0.5$, taking advantage that CSs typically have high coherence and are not significantly affected by temporal decorrelation [21], [29]. This is illustrated in Fig. 4 with the histograms of the coherence values for the CSs (blue line) and for all the pixels (orange line) inside a very large image patch showing the city center. This comparison is done for two different temporal baselines: one of 22 days, shown Fig. 4(a), and one of almost three years, shown in Fig. 4(b). As expected, the histograms for the CSs show clear maxima for values very close to 1. This indicates that temporal decorrelation is not significant for CSs even after a period of almost three years. There are a few CSs with low coherence values, but these are most likely due to changes between the two images and also some false positives in the CS detection. The number of CSs with low coherence increases for the longer temporal baseline, as many more changes occurred during this time. When comparing the two histograms for all the image pixels, it is clear that the temporal decorrelation is in this case much more significant. Also, when considering all the image pixels, the coherence values are much more evenly distributed even for short temporal baselines. Because of all this, a fixed coherence threshold like the one used in this article works well for CSs, but it is unlikely to work well when applied to all the image pixels. When considering all the image pixels, a different CD metric such as GLRT [13] would likely perform better.

C. CD With an Image Pair Using CSs

We have claimed that the proposed method is able to detect only those changes corresponding to MMOs, and that it is not affected by temporal decorrelation. To illustrate this, we apply the method described in Section II-B to an image pair acquired over the Munich area of the “Deutsches Museum.” The first image was acquired on March 28, 2016 and the second on March 13, 2018. The two amplitude images can be seen in Fig. 5(a)

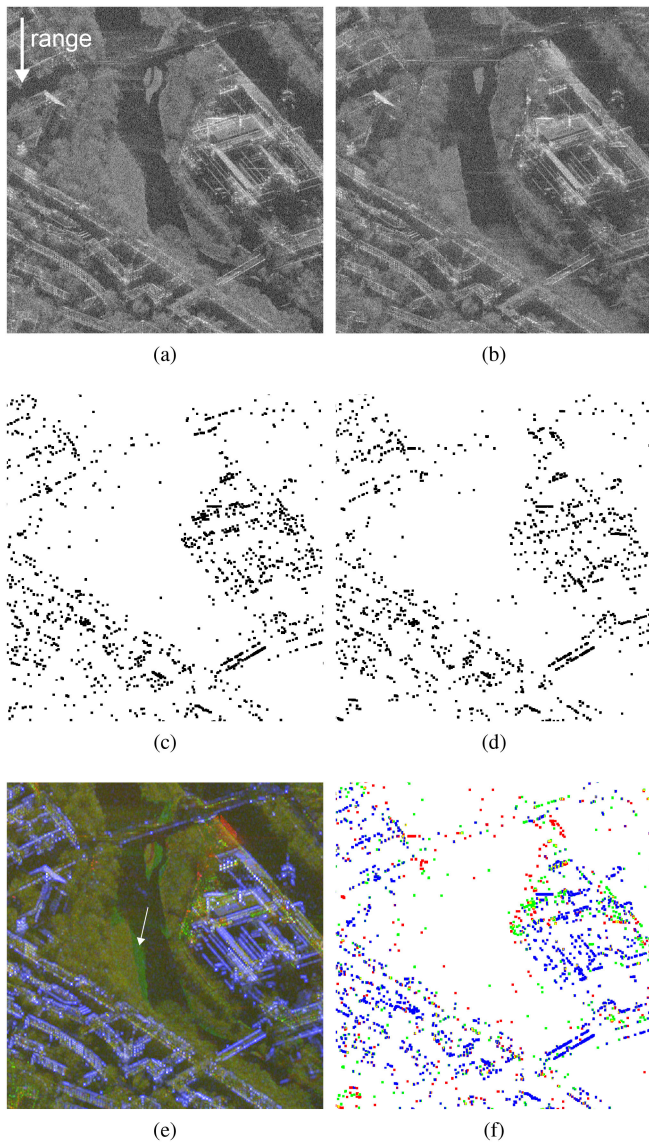


Fig. 5. CD with CSs for an image pair acquired over the Munich area of the “Deutsches Museum.” (a) SAR image from 2016. (b) SAR image from 2018. (c) Detected CSs for the 2016 image. (d) Detected CSs for the 2018 image. (e) RGB composite image with the 2018 image in red, the 2016 image in green, and the coherence in blue. (f) CSs color coded according to the type of change; blue: unchanged CS, green: CS only present in 2016, red: CS only present in 2018, and yellow: a different CS in each date.

and (b), and the detected CSs for each image are shown in Fig. 5(c) and (d), respectively. The CSs are represented as large points for better visualization, but each CS actually corresponds to an individual pixel in the full resolution SAR images. The region shown in these images contains several buildings and two bridges, as well as some vegetation and a river. As expected, the CSs are detected in the image regions where MMOs are located, with very few CSs being detected in the areas with water and vegetation. A multitemporal color composite image highlighting the changes between this image pair is shown in Fig. 5(e), with both the amplitude images in the green and red channels, and the coherence in the blue channel. In such a composite image, unchanged areas appear in blue and white, as they

exhibit low-amplitude change and high coherence. Changed areas exhibiting strong amplitude variations appear in bright green (if the amplitude decreased) and red (if it increased). Changes due to a loss of coherence with no significant amplitude variation (e.g., due to temporal decorrelation) appear in brownish and yellow colors. After applying the proposed CD method, the CSs detected in the image pair are classified into the different types of change listed in Table I. The results are shown in Fig. 5(f) with the CSs color coded according to the change type. For consistency, the colors were chosen to be similar to those in the color composite image shown in Fig. 5(e). Unchanged CSs are shown in blue, those that were only present on the first or second images are shown in green and red, respectively, and pixels containing a different CS in each image (or one that changed) are shown in yellow.

This example shows that the proposed CD approach can successfully detect the changes associated with MMO (e.g., in the “Deutsches Museum,” the large building toward the right part of the image). Changes associated with natural targets, like the change in water level causing a strong amplitude variation at the river bank [signaled by a white arrow in Fig. 5(e)], are mostly ignored. As expected, CSs are not affected by temporal decorrelation even with a temporal baseline of approximately two years. All the CSs in the unchanged buildings and other objects exhibit high coherence and are correctly detected as unchanged. Finally, a few isolated CSs that appear to be false detections (e.g., those in the river) can also be seen. While these are very few, they could be avoided by decreasing the threshold T of the CS detection. However, this would result in an increased number of false negatives (i.e., undetected CSs). Instead, these few false detections are handled during the clustering and segmentation steps.

D. Coherence Matrix of Transient Changes

In Section II-C, we introduced the concept of transient changes and briefly explained how they can be identified by their characteristic coherence matrix. To illustrate this, a real example of a transient change due to a building’s roof temporarily covered with snow can be seen in Fig. 6. A multitemporal color composite image comparing the first and last images of a time series with eight images is shown in Fig. 6(a). These two images were acquired almost two years apart. Some changes (highlighted in bright red and green) are visible toward the top, but the circular building in the center remains unchanged, as shown by its blue and white colors indicating a high coherence. However, in the composite image comparing the second and third images (acquired approximately one month apart) shown in Fig. 6(b), it appears that this same circular building has changed. In this case, the coherence of the corresponding pixels is low. The visual interpretation of the full extent of the imaged scene for this third image, acquired during wintertime, suggests that snow is the reason for this low coherence. This same effect can be seen across many other buildings over the whole city, and changes in backscatter consistent with snow cover can also be seen at many other locations (e.g., at the sides of the streets). Fig. 6(c) shows the coherence matrix for one of the CSs of the building [the one

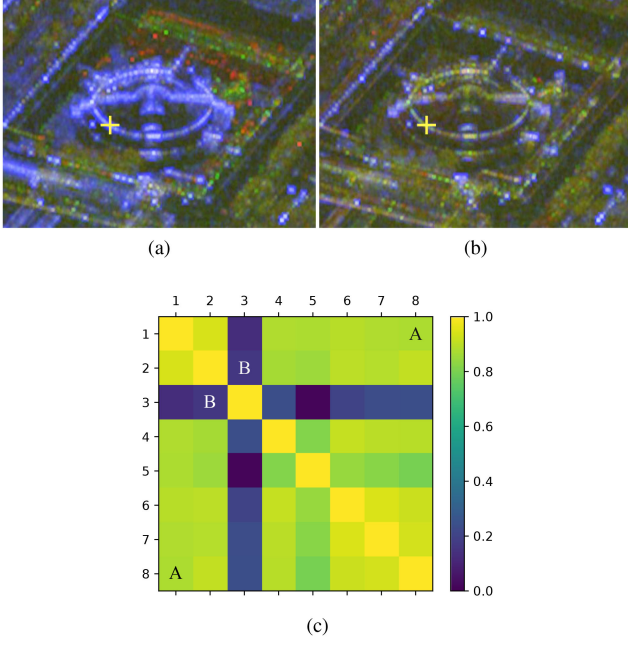


Fig. 6. Example of a transient change: circular building covered with snow in the third image of a time series. Color composite images comparing the (a) first and last images and (b) the second and third images. (c) Coherence matrix with the coherence values for all the image pairs for the highlighted pixel. The matrix elements for the pairs in (a) and (b) are marked with “A” and “B,” respectively.

marked with a yellow cross on Fig. 6(a) and (b)]. The matrix elements contain the coherence values for different image pairs: the values for the pairs shown in Fig. 6(a) (images 1 and 8) and (b) (images 2 and 3) are marked with an “A” and “B,” respectively. For this particular example, the third image is very clearly an outlier, as all the image pairs containing this image (i.e., third column and third row) have very low coherence, whereas all the other image pairs have high coherence values.

For comparison, the change caused by the construction of a building occurring at a different location in the same time series can be seen in Fig. 7. The visual analysis of the time series shows that the construction work finished between the acquisition of the fourth and fifth images. These two images are compared in the color composite image of Fig. 7(a). Strong amplitude changes highlighted in bright green and red colors can be seen across the complete façade of the building at the center of this image. Toward the top, a construction crane can also be seen in green, indicating that it was only present in the fourth image. After the fifth image, the newly constructed building remains unchanged. This can be seen in Fig. 7(b): a composite image comparing the fifth and last images. This temporal behavior can be clearly seen in the coherence matrix, shown in Fig. 7(c) for the pixel highlighted with a yellow cross in Fig. 7(a) and (b). As the building construction was finished between t_4 and t_5 , the coherence $\gamma_{j,k}$ for all the image pairs with $j \leq 4$ and $k \geq 5$ is low. This behavior is clearly different from the one exhibited by the previously shown transient change. In addition, the fact that this building then remains unchanged after t_5 can also be seen in this matrix, as $\gamma_{j,k}$ is high for all the image pairs with $j \geq 5$ and $k \geq 5$.

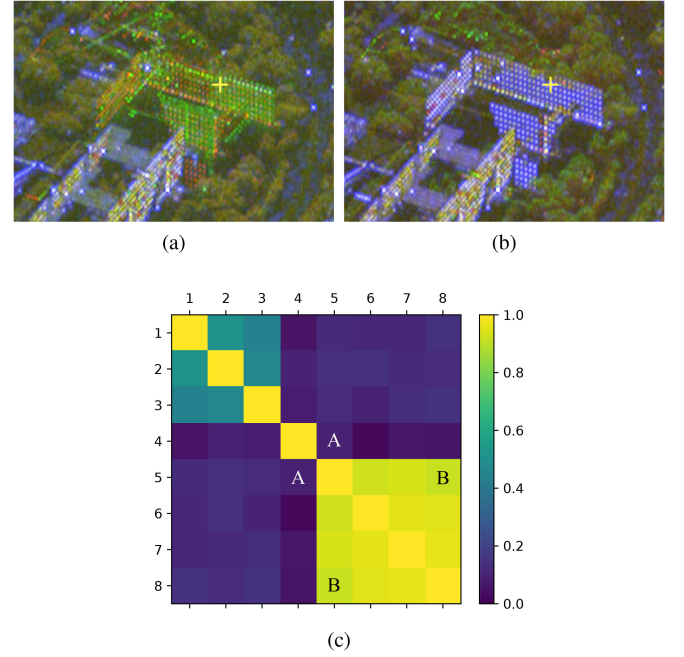


Fig. 7. Example of a change due to construction work finished between the fourth and fifth images of a time series. Color composite images comparing the (a) fourth and fifth images and (b) fifth and last images. (c) Coherence matrix with the coherence values for all image pairs for the highlighted pixel. The matrix elements for the pairs in (a) and (b) are marked with “A” and “B,” respectively.

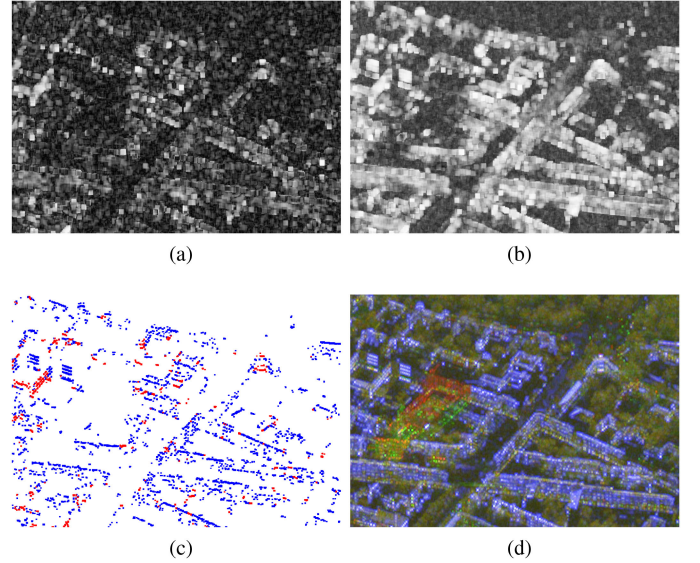


Fig. 8. Example of the effect of r in the CD metric defined in (2) for a time series of eight images. Metric f_2 computed (a) with $r = 0$ and (b) with $r = 2$. (c) Unchanged CSs detected using these two metrics: red CSs detected with both, blue only with $r = 2$. (d) As reference, a color composite image comparing the first and last images.

E. CD With a Time Series Using CSs

Using a time series instead of an image pair introduces two additional parameters. The first one is r , related to the CD metric defined in (2). The effect of r is illustrated with an example in Fig. 8, using the same eight images as in Figs. 6 and 7. The resulting CD metric f_2 for detecting changes between images 2

and 3 of this series is shown in Fig. 8(a) and (b) for $r = 0$ and $r = 2$, respectively. The metric computed with $r = 0$ has low values (indicating change) even for unchanged buildings, whereas the metric computed with $r = 2$ correctly has high values there. This difference between both the metrics is due to the previously described transient changes, caused by snow in this example. To further illustrate this, Fig. 8(c) shows a comparison of the detected CSs that remain unchanged along the complete time series according to both the metrics (i.e., those with $\min_i f_i \geq \gamma_t$). The CSs highlighted in red are detected as unchanged using both $r = 0$ and $r = 2$, whereas those in blue only with $r = 2$. The blue CSs were, therefore, affected by a transient change at some point during the time series. For visual verification, a color composite image comparing the first and last images of this series and showing the unchanged buildings in blue can be seen in Fig. 8(d). This example shows that higher values of r increase robustness against transient changes, as expected. When processing the complete time series with 49 images, we set $r = 5$ for even more robustness against transient changes, at the cost of slightly increased computation time.

The second parameter is k , related to the postprocessing step for discarding CSs that are likely false positives. Our experiments with TerraSAR-X data showed that a value of $k = 0.1$ appears to work well for time series of different lengths. Higher values of k result in more CSs being discarded.

F. Spatiotemporal Clustering of CSs

The values for the parameters of the clustering step depend on the amount and density of CSs in the objects to be detected. These two factors mainly depend on the resolution of the input SLC images and the object size. Higher resolution typically results in an increased number of detected CSs. As mentioned in Section II-A, the image with the detected CSs has virtually the same resolution as the input image. The lower resolution of the coherence images used for the CD should not play a significant role for the clustering, as it does not affect the number of detected CSs.

For all the examples shown in this article, the following parameters were used for the DBSCAN algorithm: $\varepsilon = 15$ m and $p = 20$. In this article, to filter out changes too small to be of interest, clusters containing less than 30 CSs or with a convex hull area smaller than 20 m^2 are discarded. As described in Section II-D, distances and areas are computed in meters for the clustering, to compensate the different pixel spacings along range and azimuth. For the data used in this article, a pixel represents a ground area of around 75×17 cm.

An example of the application of this clustering step can be seen in Fig. 9. Three small SAR image patches [see Fig. 9(a)–(c)] show the construction of a building from start to finish. These are part of a time series with eight images, which were processed as described in Section II-C. As the goal of this example is to illustrate the clustering step in a simple way, only a subset of the detected changing CSs is shown: those appearing in the sixth image and still present in the last image. Fig. 9(d) shows these CSs highlighted in red over the last SAR image of the series. Most of these correspond to the newly constructed building,

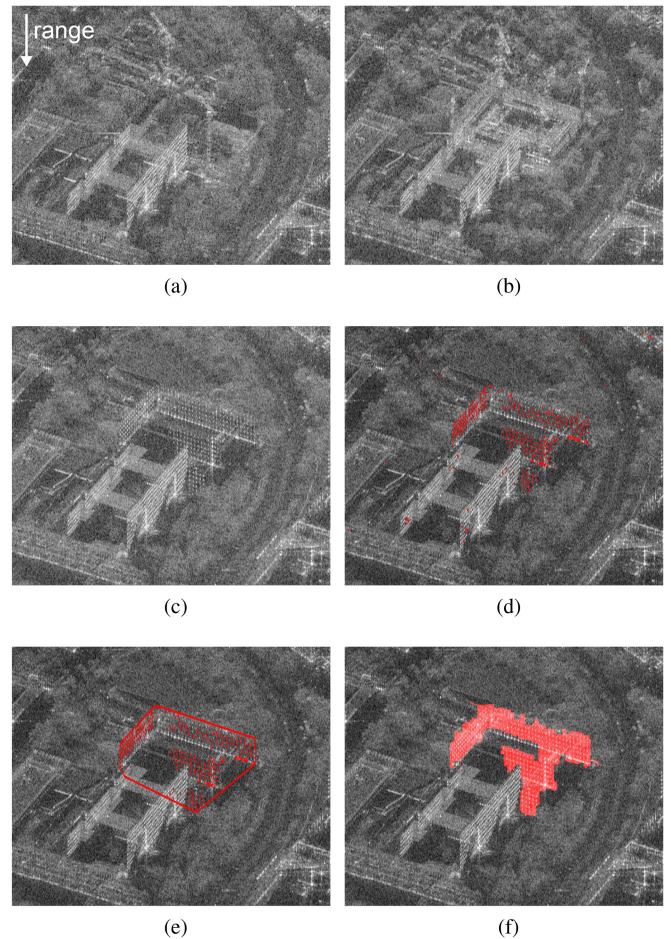


Fig. 9. Example of object-based CD. (a)–(c) Three of the SAR images in a time series with eight images showing the construction of a building. (d) CSs in the last image, which first appeared in the sixth image. (e) Resulting cluster delimited by its convex hull. (f) Segmented change.

which was finished sometime in between the acquisition of the fifth and sixth images. The clustering results obtained with the parameters listed above can be seen in Fig. 9(e). This example illustrates how the proposed spatiotemporal clustering can successfully group together all the CSs belonging to a changed object.

G. Segmentation of the Detected Changes

The proposed method for change segmentation has a few parameters only. The first parameter is a scaling factor, used to compute the extent of the image patches to be considered for the segmentation. In this article, we set this factor to 50%, meaning that these patches are 50% larger than the rectangles enclosing the corresponding clusters. Another parameter is the radius for the closing operation performed during the postprocessing of the obtained segmentation mask. We set this to 5 pixels. The selected values seem to work well for the applications considered in this article. An example of the segmentation results obtained for a change due to the construction of a new building is shown in Fig. 9. The cluster with the change to be segmented and the resulting segmentation are highlighted in red in Fig. 9(e) and (f),

respectively. Using the temporal information obtained from the CD and the clustering results as a starting point, the proposed method achieved a rather accurate segmentation of the newly constructed building.

Additional parameters are needed for the segmentation of objects present in only one image, as the SAR amplitude must also be considered for this. The first one is a threshold Δ_A , applied to the amplitude difference in dB scale. We set this threshold to $\Delta_A = 3$ dB. The second one is a threshold A_{\min} related to the minimum amplitude value for pixels corresponding to MMOs, as these tend to exhibit relatively high-amplitude values. This one is set to $A_{\min} = -15$ dB. Before applying these thresholds, the speckle noise should be reduced. In this article, we apply a custom speckle filter that preserves the full resolution. Multilooking is applied to the despeckled image to further reduce speckle, resulting in a 1-m resolution in slant range and azimuth, but keeping the original pixel spacing. We do not describe this speckle filter in detail as this is out of the scope of this article. However, modern despeckling algorithms, such as [33] and [34], would very likely result in a better speckle reduction. Our experiments show that these two fixed thresholds work quite well for our data, as they are only applied locally where changes have already been detected for many CSs. Some example results can be seen later in Section IV-B.

IV. RESULTS

The CD method proposed in this article was applied to the complete time series of the city of Munich introduced in the previous section. The processing chain was applied as described in Section II and illustrated in Fig. 1. The different parameters were set to the values suggested in Section III. These values were empirically chosen and have been shown to work well in a variety of settings. Many different types of changes can be seen in the imaged scene during this period of almost three years. In this article, we focus on the changes due to construction activity, such as newly constructed or renovated buildings and infrastructure, the build-up for festivals, etc. To illustrate the capabilities of the proposed method to categorize the detected changes, we show how it can be applied to specifically detect some of these types of changes.

A. Detection of New and Renovated Buildings

The detection of changes due to the construction of new buildings and infrastructure or renovations to existing ones is of interest for many applications. Often, there are many different changes continuously occurring across the whole imaged scene, and most general CD methods simply result in a binary change map highlighting all the changes. In contrast, the method proposed in this article can identify these changes by their characteristic temporal behavior, as described in Section II-F. This approach identifies the final change to a given building and the date when it happened, which should correspond to time when the construction work was finished. For this example, ΔT was set to two months.

The obtained results for an area around the city campus of the Technical University of Munich (TUM) are shown in

Fig. 10(a). This image shows the latest SAR image in the series (acquired on February 28, 2019), with the segmented changes highlighted in different colors according to the date when the construction work finished. The visual inspection of the SAR time series allows us to validate the obtained results, as it can be seen that renovation works have actually taken place at all the highlighted locations. In some cases, these renovations resulted in visible alterations to the corresponding buildings, while in others they could only be identified by the temporary presence of scaffolds. Many other changes occurred in this area and were also detected by the proposed method, but all those not fitting the desired temporal pattern were discarded. This successfully discarded changes corresponding to moving objects and other activity. However, it also discarded two changes due to the construction of two new buildings, which were still not completely finished by the acquisition time of the last image.

Further verification using optical images has been performed for the museum “Alte Pinakothek” and for one side of the TUM building, both areas enclosed in Fig. 10(a) by white rectangles. The renovation process of the “Alte Pinakothek” can be seen in three optical images at the bottom of Fig. 10(l) and (m) and the five SAR images in the second row of Fig. 10(b)–(f), each with their corresponding acquisition date. In the same way, the changes at the TUM building can be seen in an optical image and a street level photography in Fig. 10(o) and (p), respectively, as well as in the five SAR images in the third row of Fig. 10(g)–(k). In both the cases, the results shown in Fig. 10(a) agree with the sequence of the renovation process depicted in both the optical and SAR images, where proof of these changes was manually outlined using the same colors. The optical images show all the renovation works at both buildings, except for the one highlighted in green on the university building in Fig. 10(a), as no optical images acquired at that time could be found. The SAR image sequences do, however, show all the detected changes, including that one: the scaffold can be seen in Fig. 10(g) and no longer in Fig. 10(h) (in both the cases highlighted in green). In some cases, the segmented changes shown in Fig. 10(a) are smaller than the manually outlined areas in the SAR and optical images. The reason for this is that the results shown in this figure only highlight the parts of the final building structures that were modified, and not all the image pixels which changed at some point throughout the time series (e.g., those temporarily covered by scaffolds). Also, the renovation of the building on the right side of the street in Fig. 10(p) could not be detected, as this façade is not visible to the SAR sensor and appears as a shadow area in the SAR images. This change could also be detected by applying the proposed method to a time series acquired with descending orbit.

In addition to the results shown for the TUM area, many more changes due to the construction or renovation of buildings were detected at different locations across the city. Fig. 11 shows four of the detected changes, each displayed in a different row [rows (a)–(d)]. Each change is illustrated by three SAR image chips acquired before (column A), during (column B), and after (column C) the construction work. The detected changes are highlighted in the final images (i.e., column C in Fig. 11),

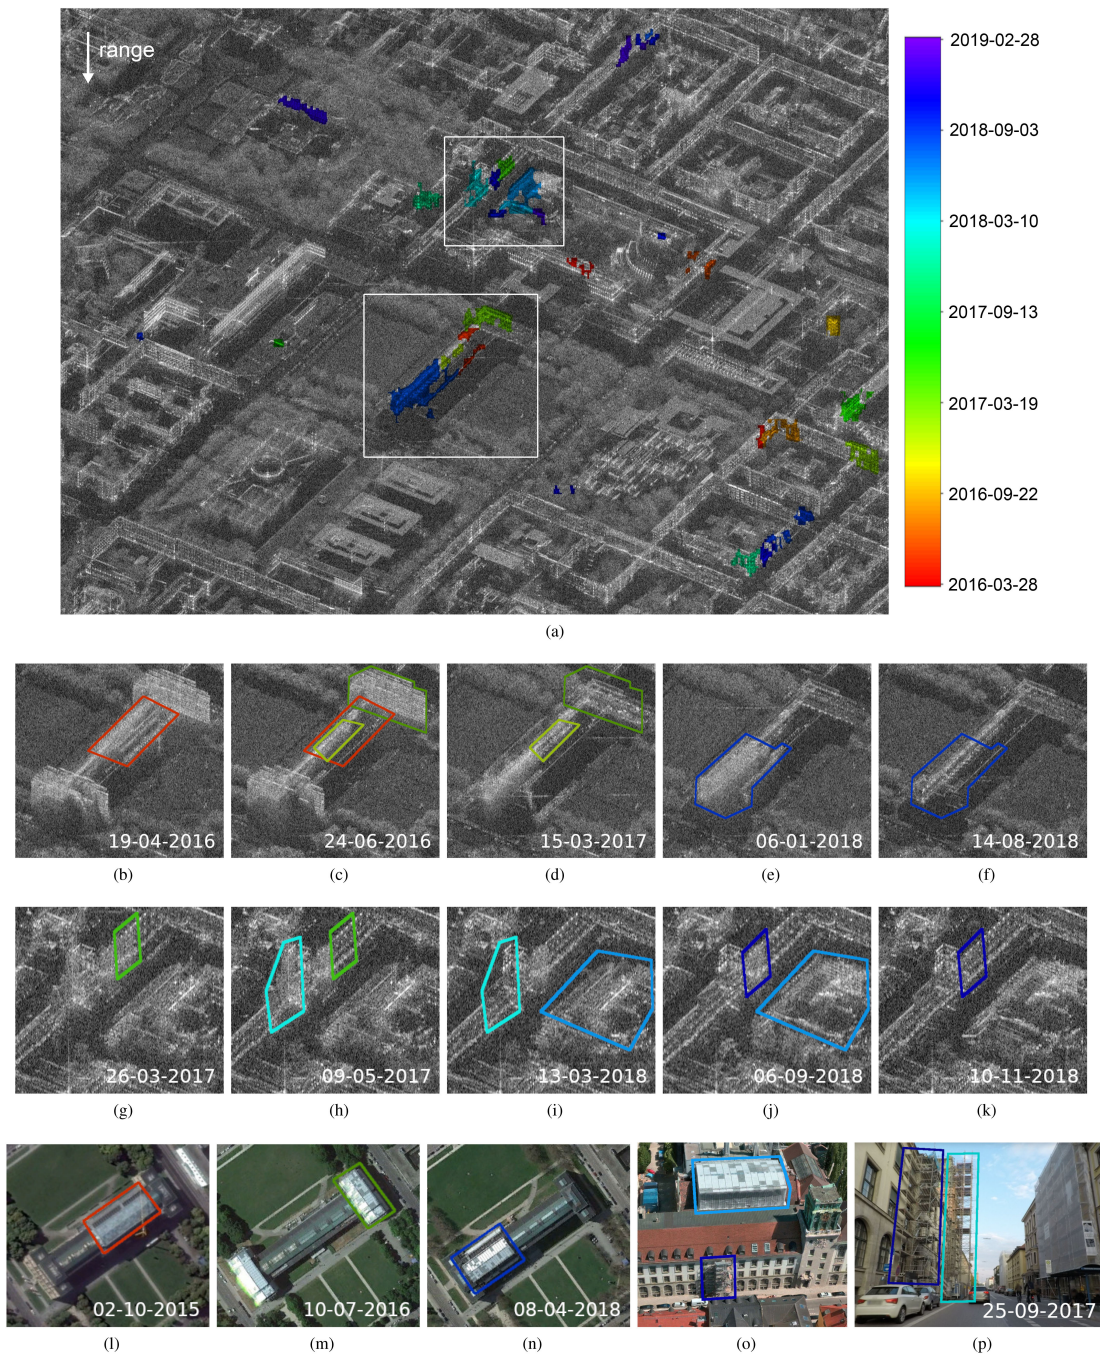


Fig. 10. Results for the detection of changes due to the construction and renovation of buildings in an area around the city campus of the TUM. For visual verification, some optical images and a few of the SAR images in the series are shown for two buildings: the museum “Alte Pinakothek” and a side of the university building. (a) Segmented changes drawn over the last SAR image of the series, with the color denoting the date at which the construction work was finished. (b)–(f) Sequence of SAR images for the museum. (g)–(k) Sequence of SAR images for the university. (l)–(n) Optical images for the museum. (o) and (p) Optical images for the university. The optical satellite images were obtained from Google Earth (l–n) and Apple Maps (o), and the street level photography from Mapillary (p). The acquisition date of the optical image in (o) was not listed in Apple Maps, but it appears to have been acquired in mid-2018.

with the colors corresponding to the date on which the construction work finished, using the same legend from Fig. 10(a). Fig. 11(a) and (b) shows changes due to the construction of new buildings, whereas Fig. 11(c) and (d) shows examples due to renovations. Some of these changes appear in multiple colors, meaning that different structures were finished at different times [like in Fig. 11(b)], or that different sections of the scaffold

were gradually removed [like in Fig. 11(d)]. Again, the visual inspection of the SAR images in the time series allowed to validate the accuracy of the obtained results. For the two new buildings, the image sequence shows an empty lot, followed by the construction progressing and finally the finished building. On the other hand, for the renovations, the first and final states are very similar, with the presence of a scaffold during the renovation

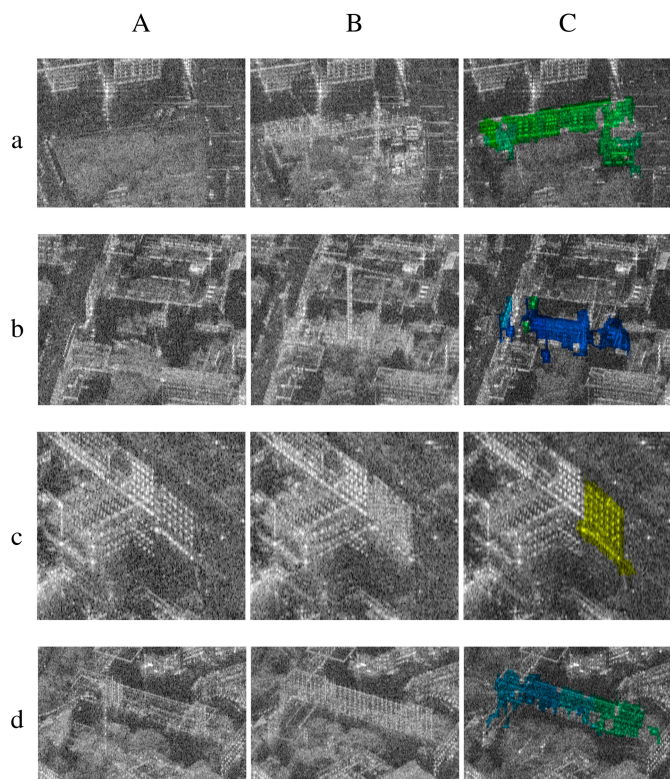


Fig. 11. Four examples [rows (a)–(d)] of detected changes to buildings. (A) Images acquired before the construction work. (B) Images acquired during the construction work. (C) Final images with the detected new and/or modified structures highlighted.

being the main change. Even if these two types of changes are different, the proposed method cannot distinguish them as they exhibit similar temporal patterns (i.e., the appearance of new CSs that later remain static).

One possible way to distinguish changes due to the construction of new buildings and the renovation of existing ones would be to evaluate the number of CSs over time inside the segmented changes. It has been shown that a significant increase in the number of CS indicates an increase in the amount of built-up structures [35]. This information is plotted in Fig. 12 for the four changes shown in Fig. 11(a)–(d). As expected, for the examples corresponding to the construction of new buildings [see Fig. 11(a) and (b)], the number of CSs increases by a significant amount as the construction progresses and remains constant after it finishes. For the renovations [see Fig. 11(c) and (d)], there is only a small variation in the number of CSs between the prior and final states, with a larger variation during the time where the scaffold is present. This kind of analysis could potentially be used to achieve a finer classification of the detected changes. However, this task is out of the scope of this article.

B. Detection of Other Changes

In addition to the construction and renovation of buildings and infrastructure, the proposed method also detects other changes like the build-up for festivals and/or any events taking place, the

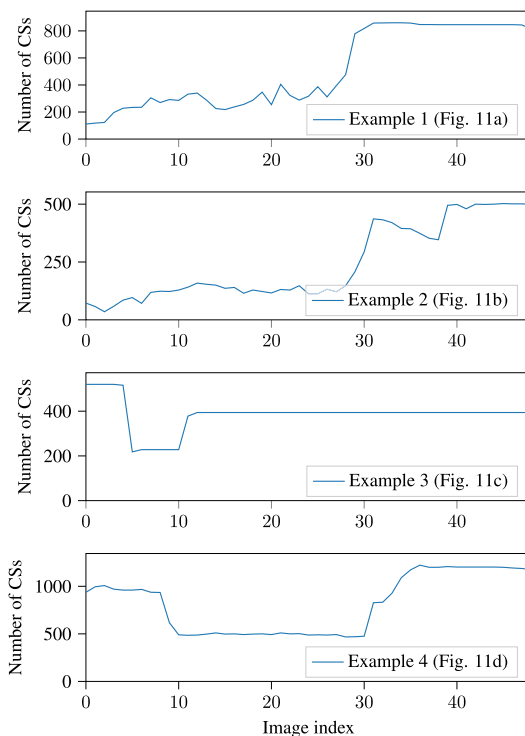


Fig. 12. Evolution of the number of CSs over time for the four changes in Fig. 11, with those corresponding to the construction of new buildings [see Fig. 11(a) and (b)] and renovations [see Fig. 11(c) and (d)] showing distinct patterns.

movement of objects, etc. As mentioned in Section II-F, such changes typically involve CSs both appearing and disappearing inside relatively short periods. To show this, the proposed approach was applied to detect only changed objects that remain unchanged and static for less than an interval ΔT of four months. Some of the results are shown in Fig. 13 for a small area around the park “Theresienwiese,” where different festivals take place throughout the year. Fig. 13(a)–(c) shows this area at three different times along with the corresponding changes. The detected changes were highlighted in a color ranging from red to yellow (with the hue component varying linearly from 0 to 60) according to the length of time that each object is present and remains unchanged. Most of these changes appear in dark red (hue value of 0), meaning that the corresponding objects were only present during one image. A few objects are highlighted in orange, such as the roofs of some of the large festival tents in Fig. 13(c), showing that these structures were built before and/or stayed longer in the scene than the other objects.

In addition to the changes due to the festivals in this park, changes were also detected for some buildings where construction work is taking place. The analysis described in the previous subsection allowed to detect the final change to buildings and estimate when the construction or renovation finished. However, during the construction phase, the buildings change continuously, and all these fast changes are also detected. Once the final change to a building is detected, all the previous changes in the same overlapping area could be identified as the corresponding construction work.

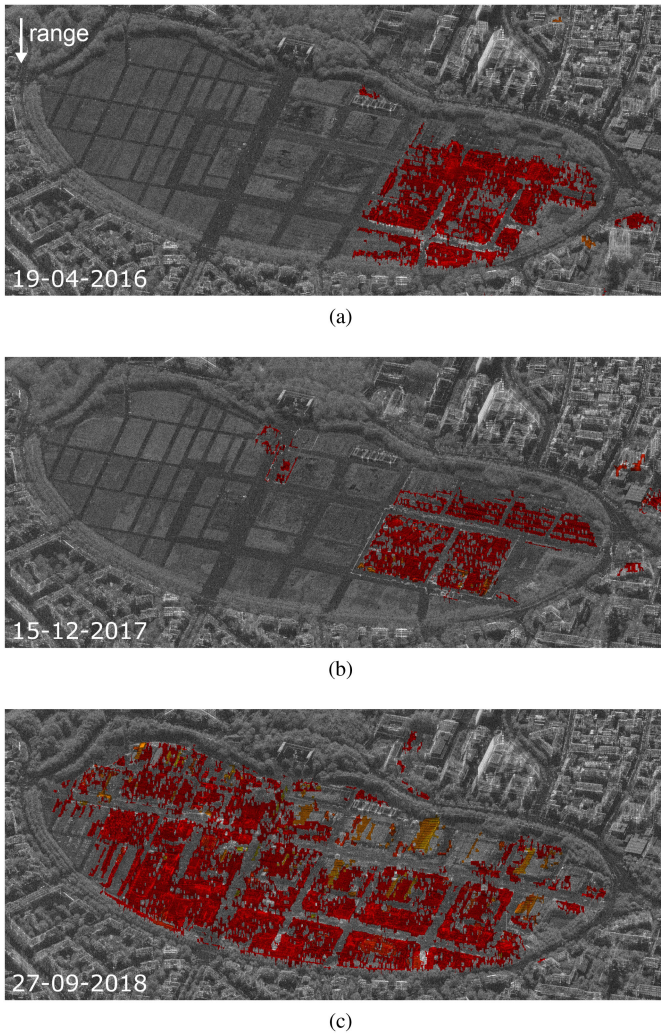


Fig. 13. Detected changes at the park “Theresienwiese” due to objects both appearing and disappearing within a time period of less than four months, highlighted in a color according to the length of time that each object is present. (a)–(c) show this park at three different dates, with different events taking place, which cause the majority of the detected changes.

C. Detection of Unchanged and Static Objects

The proposed method also detects the MMOs that remain static and unchanged throughout the time series. For an urban scene like the one in our dataset, these correspond to the existing buildings and infrastructure where no renovations took place. Fig. 14 shows an example of the obtained results for an area close to the Munich city center. The unchanged CSs are highlighted in blue over the first and last images of the time series in Fig. 14(a) and (b), respectively. These CSs are similar to the persistent scatterers of PSI methods [21], but can be obtained with as few as two images. However, especially in layover areas (e.g., the building façades), the resolution loss in range due to the sublookings required for the CS detection results in less point scatterers being detected than when using PSI. The proposed object-based analysis method was also applied to segment the unchanged objects. The segmentation results can be seen in Fig. 14(c) and (d), again highlighted over the first and last images




Fig. 14. Objects that remained unchanged throughout the time series. Top: detected unchanged CSs highlighted in blue over (a) the first image and (b) the last image. Bottom: segmentation results for the unchanged objects highlighted in blue over (c) the first image and (d) the last image.

of the series, respectively. The comparison of the left and right columns (corresponding to the first and last images) of Fig. 14 show how the proposed method correctly identifies unchanged objects, consisting mostly on buildings, but also street lamps, etc. For some of the buildings that are not highlighted, the changes can be clearly seen (e.g., one building toward the top, and another one toward the bottom left). For others, no changes can be seen between the first and last images, but the visual inspection of the complete time series revealed that renovations took place.

V. DISCUSSION

We introduced a novel method for the detection of changes associated with MMO using pairs or series of SAR images. This method was applied to a time series of 49 TerraSAR-X images for the monitoring of construction activity. The analysis of the results shows that the method performs well and can accurately detect these changes. Seasonal changes (e.g., snow, changes in water level, vegetation, etc.) are ignored and do not result in false alarms.

When applying the method for detecting changes due to the renovation of existing buildings, the combined use of multiple time series acquired with different imaging geometries (e.g., different orbit or look direction) should be considered. Otherwise, some changes cannot be detected, as all the building façades cannot be imaged with a single imaging geometry.

For the build-up of festivals and similar events involving many small objects close together and constantly changing, the proposed method is not able to separate individual objects using neither spatial nor temporal information. In such cases, closely packed objects are grouped together when performing the object-based analysis. Other than that, the change segmentation works fairly well.

The proposed CD method was evaluated in a rather qualitative way, as quantitatively evaluating its performance and comparing it to other CD methods would require a dataset with ground truth. To the best of our knowledge, there is no such publicly available dataset, and generating one is not an easy task. Acquiring ground truth or manually labeling the data for such a CD task is very challenging and time consuming. Besides, time series of spaceborne SAR data with such a high resolution are typically not freely available. Also, the generation of realistic synthetic data does not seem feasible, as the proposed method uses SLC SAR images and exploits their phase.

The presented method was developed for a specific task: detecting changes associated with MMOs. For this task, we expect it to perform better than general-purpose CD methods. However, this specificity limits its applicability: it is not well suited for applications where changes to natural targets are relevant. Also, it will perform badly when applied to low-resolution data (e.g., Sentinel-1), as the CS detection requires a large bandwidth to work properly. On the other hand, better performance can be expected when using data with even higher resolution. The proposed method can work with long temporal baselines, but the potential of the applied temporal analysis increases with the temporal resolution of the used time series. This makes it especially interesting for SAR missions involving large constellations with very high revisit, like those currently being built by New Space companies [36].

VI. FUTURE WORK

Further work will involve applying the proposed method to datasets with different scenes to detect other kinds of changes, such as the arrival and departure of airplanes at airports and ships at ports. In addition, we will also explore the possibility of exploiting multiple images jointly for CS detection, instead of performing the detection separately in each image and then applying a consistency check as a postprocessing step. The segmentation of objects present in a single image of the series could also potentially be improved by using a more modern amplitude CD metric and/or a more advanced segmentation method. Finally, a more sophisticated analysis of the segmented objects could be implemented to better distinguish different kinds of changes. For example, changes due to new buildings could potentially be distinguished from those due to renovations by analyzing the evolution over time of the number of CSs, as we have briefly shown in this article.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their constructive comments and recommendations.

REFERENCES

- [1] A. Moreira, P. Prats-Iraola, M. Younis, G. Krieger, I. Hajnsek, and K. P. Papathanassiou, "A tutorial on synthetic aperture radar," *IEEE Geosci. Remote Sens. Mag.*, vol. 1, no. 1, pp. 6–43, Mar. 2013.
- [2] C. V. Lopez and U. Stilla, "Object-based SAR change detection for security and surveillance applications using density based clustering," in *Proc. 12th Eur. Conf. Synthetic Aperture Radar*, 2018, pp. 1–6.
- [3] C. V. Lopez, T. Kempf, R. Speck, H. Anglberger, and U. Stilla, "Automatic change detection using very high-resolution SAR images and prior knowledge about the scene," *Proc. SPIE*, vol. 10188, 2017, Art. no. 1018805.
- [4] C. Marin, F. Bovolo, and L. Bruzzone, "Building change detection in multitemporal very high resolution SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 5, pp. 2664–2682, May 2015.
- [5] F. Bovolo, C. Marin, and L. Bruzzone, "A hierarchical approach to change detection in very high resolution SAR images for surveillance applications," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 4, pp. 2042–2054, Apr. 2013.
- [6] C. V. Lopez and U. Stilla, "Monitoring of oil tank filling with spaceborne SAR using coherent scatterers," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 5638–5655, 2021.
- [7] M. Preiss, D. A. Gray, and N. J. S. Stacy, "Detecting scene changes using synthetic aperture radar interferometry," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 8, pp. 2041–2054, Aug. 2006.
- [8] E. J. M. Rignot and J. J. van Zyl, "Change detection techniques for ERS-1 SAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 31, no. 4, pp. 896–906, Jul. 1993.
- [9] M. Cha, R. D. Phillips, P. J. Wolfe, and C. D. Richmond, "Two-stage change detection for synthetic aperture radar," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 12, pp. 6547–6560, Dec. 2015.
- [10] A. Parizzi, X. Cong, and M. Eineder, "First results from multifrequency interferometry. A comparison of different decorrelation time constants at L, C, and X Band," in *Proc. ESA Fringe Workshop*, 2009, pp. 1–5.
- [11] A. Reigber, M. Jager, and E. Krogager, "Polarimetric SAR change detection in multiple frequency bands for environmental monitoring in arctic regions," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2016, pp. 5702–5705.
- [12] J. Tao and S. Auer, "Simulation-based building change detection from multiangle SAR images and digital surface models," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 9, no. 8, pp. 3777–3791, Aug. 2016.
- [13] A. V. Monti-Guarnieri, M. A. Brovelli, M. Manzoni, M. M. d'Alessandro, M. E. Molinari, and D. Oxoli, "Coherent change detection for multipass SAR," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 11, pp. 6811–6822, Nov. 2018.
- [14] E. C. Koeniguer and J.-M. Nicolas, "Change detection based on the coefficient of variation in SAR time-series of urban areas," *Remote Sens.*, vol. 12, no. 13, 2020, Art. no. 2089.
- [15] G. Quin, B. Pinel-Puysegur, J.-M. Nicolas, and P. Loreaux, "MIMOSA: An automatic change detection method for SAR time series," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 9, pp. 5349–5363, Sep. 2014.
- [16] K. Conradsen, A. A. Nielsen, J. Schou, and H. Skriver, "A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 1, pp. 4–19, Jan. 2003.
- [17] M. J. Canty, A. A. Nielsen, K. Conradsen, and H. Skriver, "Statistical analysis of changes in Sentinel-1 time series on the Google earth engine," *Remote Sens.*, vol. 12, no. 1, 2020, Art. no. 46.
- [18] Y. Li, C. Peng, Y. Chen, L. Jiao, L. Zhou, and R. Shang, "A deep learning method for change detection in synthetic aperture radar images," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 8, pp. 5751–5763, Aug. 2019.
- [19] 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. 2019.
- [20] J. Geng, X. Ma, X. Zhou, and H. Wang, "Saliency-guided deep neural networks for SAR image change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 10, pp. 7365–7377, Oct. 2019.
- [21] A. Ferretti, C. Prati, and F. Rocca, "Permanent scatterers in SAR interferometry," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 1, pp. 8–20, Jan. 2001.
- [22] C.-H. Yang and U. Soergel, "Adaptive 4D change detection based on PSI," in *Proc. IEEE 12th Eur. Conf. Synthetic Aperture Radar*, 2018, pp. 1–5.
- [23] S. Lobry, F. Tupin, and L. Denis, "A decomposition model for scatterers change detection in multi-temporal series of SAR images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2016, pp. 3362–3365.

- [24] S. Lobry, L. Denis, and F. Tupin, "Multitemporal SAR image decomposition into strong scatterers, background, and speckle," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 9, no. 8, pp. 3419–3429, Aug. 2016.
- [25] M. J. Sanjuan-Ferrer, I. Hajnsek, K. P. Papathanassiou, and A. Moreira, "A new detection algorithm for coherent scatterers in SAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 11, pp. 6293–6307, Nov. 2015.
- [26] R. Z. Schneider and K. Papathanassiou, "Estimation and correction of ionospheric induced phase errors in SAR images using coherent scatterers," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2009, pp. IV-165–IV-168. [Online]. Available: <https://ieeexplore.ieee.org/document/5417348>
- [27] V. M. Giacomazzo, A. Refice, F. Bovenga, and N. Veneziani, "Identification of coherent scatterers: Spectral correlation vs. multi-chromatic phase analysis," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2008, pp. IV-411–IV-414. [Online]. Available: <https://ieeexplore.ieee.org/document/4779745>
- [28] R. Z. Schneider, K. P. Papathanassiou, I. Hajnsek, and A. Moreira, "Polarimetric and interferometric characterization of coherent scatterers in urban areas," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 4, pp. 971–984, Apr. 2006.
- [29] M. Eineder, N. Adam, R. Bamler, N. Yague-Martinez, and H. Breit, "Spaceborne spotlight SAR interferometry with TerraSAR-X," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 5, pp. 1524–1535, May 2009.
- [30] E. Sansosti, P. Berardino, M. Manunta, F. Serafino, and G. Fornaro, "Geometrical SAR image registration," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 10, pp. 2861–2870, Oct. 2006.
- [31] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "Density-based spatial clustering of applications with noise," in *Proc. Int. Conf. Knowl. Discov. Data Mining*, 1996, vol. 240, p. 6.
- [32] Y. Bazi, L. Bruzzone, and F. Melgani, "Automatic identification of the number and values of decision thresholds in the log-ratio image for change detection in SAR images," *IEEE Geosci. Remote Sens. Lett.*, vol. 3, no. 3, pp. 349–353, Jul. 2006.
- [33] C.-A. Deledalle, L. Denis, F. Tupin, A. Reigber, and M. Jager, "NL-SAR: A unified nonlocal framework for resolution-preserving (Pol)(In)SAR denoising," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 4, pp. 2021–2038, Apr. 2015.
- [34] E. Dalsasso, L. Denis, and F. Tupin, "As if by magic: Self-supervised training of deep despeckling networks with MERLIN," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–13, 2022.
- [35] C. V. Lopez and U. Stilla, "Using coherent scatterers in time series of high resolution SAR images for the monitoring of construction activity," *ISPRS Ann. Photogrammetry Remote Sens. Spatial Inf. Sci.*, vol. IV-2/W7, pp. 183–187, 2019.
- [36] D. Muff et al., "The ICEYE constellation—Some new achievements," in *Proc. IEEE Radar Conf.*, 2022, pp. 1–4.



Carlos Villamil Lopez was born in Cervo, Spain, in 1989. He received the M.Sc. degree in telecommunications engineering from the University of Oviedo, Gijón, Spain, in 2013. He is currently working toward the Ph.D. degree with the Unit "Photogrammetry and Remote Sensing," Technical University of Munich, Munich, Germany. His Ph.D. study concerns change detection with time series of high-resolution SAR images.

In 2012, he was a Visiting Student with Virginia Tech, Blacksburg, VA, USA. Since 2013, he has been a Research Scientist with the Microwaves and Radar Institute, German Aerospace Center, Wessling, Germany. His main research interests include synthetic aperture radar (SAR) image processing and SAR change detection.



Uwe Stilla (Senior Member, IEEE) was born in Cologne, Germany, in 1957. He received the Dipl.-Ing. degree in communications engineering from the University of Paderborn, Paderborn, Germany, in 1980, and the Dipl.-Ing. degree in electrical engineering and the Ph.D. degree in electrical engineering (pattern recognition) from the University of Karlsruhe, Karlsruhe, Germany, in 1987 and 1993, respectively.

From 1987 to 1990, he was with the Institute of Biomedical Engineering, University of Karlsruhe. From 1990 to 2004, he was with the FGAN-FOM Institute of Optronics and Pattern Recognition (now Fraunhofer IOSB, Ettlingen, Germany). Since 2004, he has been a Professor with the Technical University of Munich (TUM), Munich, Germany. He is the Head of the Unit "Photogrammetry and Remote Sensing" with the TUM Department of Aerospace and Geodesy and was the Dean of Studies (from 2005 to 2016) and Vice-Dean (from 2008 to 2013) of the TUM Department of Civil Engineering and Geodesy. He organized and co-organized many conferences and workshops of the International Society for Photogrammetry and Remote Sensing, the IEEE, and the International Association for Pattern Recognition, and authored or coauthored more than 500 scientific contributions. His research interests include image analysis in the field of photogrammetry and remote sensing.

Dr. Stilla is a Member of the Scientific Board of the German Commission of Geodesy and the Vice-Chairman of the Commission for Geodesy and Glaciology of the Bavarian Academy of Science and Humanities. He was the Vice-President (from 2012 to 2016) and the President (from 2016 to 2020) of the German Society for Photogrammetry, Remote Sensing and Geoinformation. Since 2004, he has been a Chair of multiple working groups of the International Society of Photogrammetry and Remote Sensing (last: Intercommission Working Group ICWG II/III "Pattern Analysis in Remote Sensing").