HOMOLOGICAL PERSISTENCE FOR SHAPE BASED CHANGE DETECTION BETWEEN DIGITAL ELEVATION MODELS

Bruno Vallet

IGN, Laboratoire MATIS
73 avenue de Paris
94165 Saint-Mandé, FRANCE
bruno.vallet@ign.fr
recherche.ign.fr/labos/matis/~vallet

Commission III/3

KEY WORDS: Change detection, Digital Elevation Model, shape, persistence, multilevel

ABSTRACT:

Digital Elevation Models provide an accurate description of scenes allowing change detection by radiometric independence measurements. The height information gives access to the changes in the 3D shape of a scene occurring between two dates. In this paper, a procedure to extract meaningful changes is proposed based on analysing this shape variation. Instead of filtering a difference map, a multilevel analysis of the difference map based on the theory of persistence is developed to make such analysis independent from a choice of threshold values.

1 INTRODUCTION

1.1 Context

Change detection from aerial or spatial images is a major stake for accelerating and automatizing of geographical databases. Such databases are becoming larger, more complete and more accurate, so the task of keeping them up to date is becoming ever more challenging. In this study, we consider change detection based solely on Digital Elevation Models (DEMs). We have made this choice because we are interested in finding only where the geometry of the scene has changed, not only its radiometry. This is much simpler as radiometric differences may come from many other sources than a change in the scene geometry: change in lighting conditions, in object color, in acquisition conditions (different viewpoint), etc. The drawback of this choice is that we will be very dependant on the quality of the DEM used to detect the changes. In particular, we will not be able to sort easily between actual changes and DEM errors.

1.2 Related works

Change detection has often been performed in order to update an existing geographic database, for instance of buildings, in which case the data at the former date is this database, and aerial/satellite imagery at the later date. This is out of the scope of this paper, which focusses on purely DEM to DEM comparison (without requiring such a database), so we refer to the detailed performance analysis of (Champion et al., 2011) on such approaches.

Most existing change detection between DEMs rely on the following simple steps:

1. Resample the two DEMs in a common geometry (if they are not already). The height information comes either from Aerial Laser Scanning (ALS) or photogrammetric dense matching using aerial images.

2. Compute the difference map between the DEMs in this common geometry. (Kuo-hsin et al., 2006) proposes to simply give this map to an operator for visual interpretation.
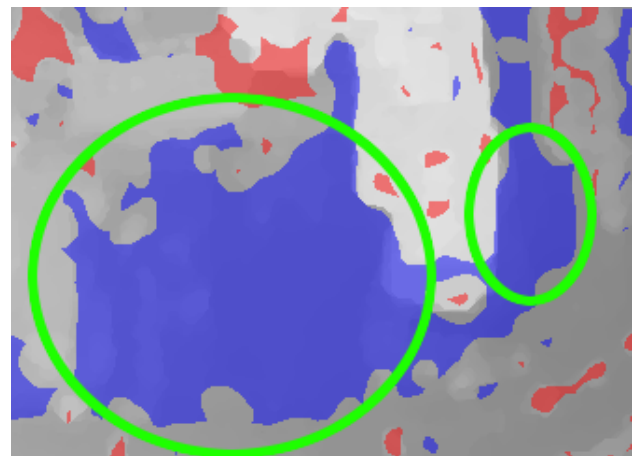

Figure 1: Constructions (blue) and destructions (red) obtained by thresholding a DEM difference map. Two separate constructions (circled in green) are merged by the choice of a low threshold. The resulting component has a low compactness while the individual real changes were highly compact, which may result in their rejection based on shape analysis.

3. Threshold the difference map

4. Filter the resulting binary mask based on the shapes of its components

The main differences rely on the way the resulting components are filtered, and sometimes the addition of other sources of information. The easiest and most popular way to filter the components is to apply some mathematical morphology filters (shrinking, expansion) which removes thin, elongated components. The radiometric information from orthophotos at the two dates can also be be used as an additional indicator of change (Sasagawa et al., 2008). A finer geometric analysis of the components can be performed through the use of a compactness factor in addition to mathematical morphology (Chaabouni-Chouayakh et al., 2010). In this case the DEMs can be classified using the Iso-Data unsupervised classification algorithm to enhance the detec-

tion. Finally the uncertainty on the height given by the individual DEMs can be estimated to refine the study of their difference in the context of topographic survey of sediment budgets (Wheaton et al., 2010). The spatial variability of elevation is then estimated through a fuzzy inference system then this estimate is modified based on the spatial coherence of height variation.

Such approaches suffer from a major drawback: The results highly depend on the choice of the threshold used to binarize the difference map. High thresholds tend to miss some significant changes, leading to high under detection rates. Conversely, low thresholds tend to merge different changes into a single change with very complex shape (Figure 1) that risks rejection by a compactness criteria while the individual changes were compact enough. The resulting detection will also be highly undersegmented.

Classification may be used in two different ways for change detection:

1. Classifying the data at the two dates and comparing the classifications, as proposed by (Chaabouni-Chouayakh et al., 2010)

2. Posing the change detection problem as a classification problem

Two very innovative recent works take the second choice:

1. (Guérin et al., 2012) poses the problem of change detection as a labelling problem with a data attachment term and a regularization term, which solution is found by dynamic programming. Thus unlike (Chaabouni-Chouayakh et al., 2010), compactness is not used to filter out components but is a prior to enforce that the components found are compact.

2. (Tian et al., 2013) proposes to merge segmentations of the DEM at the two dates and classify the resulting segments using various indicators.

Finally, we can cite a very different approach: in the case of aerial laser scanning, (Hebel et al., 2011) proposed to work directly in 3D (instead of a 2.5D DEM) by voxelizing space and using the intersection of laser rays with such voxels to decide wether thay are occupied. Changes are then located by finding conflicts between such occupancy computed for two different acquisitions using Dempster-Shafer evidence theory.

1.3 Method overview

The approach that we propose studies the morphology of changes in order to make the detection independent from the choice of an arbitrary (and hard to tune) threshold, thus to avoid the aforementioned issues. Homological persistence will be used to study changes at all the levels when significant changes happen in order to automatically choose the threshold the most appropriate to their analysis.

As with classical approaches, we start by filtering a difference map between the DEMs at the two dates and splitting this difference map (Section 2.1). Splitting separates positive and negative components of the difference maps (potential constructions and destructions), resulting in very large components, which only serves to cut the area of interest into smaller areas that will be analysed independently. This is preferable over cutting the area in rectangular blocks because this way we ensure that no change will overlap two blocks (changes will be strictly contained in
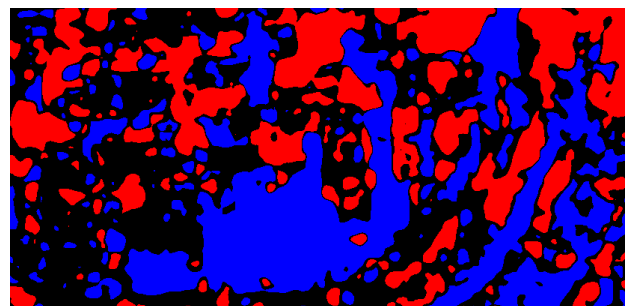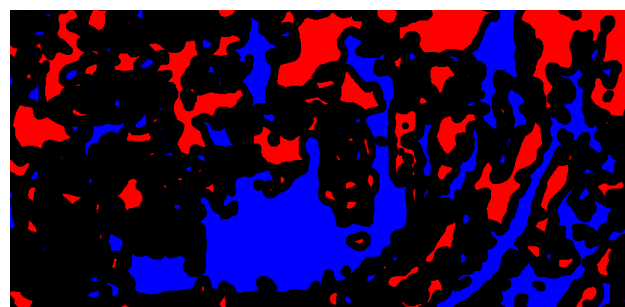

(a) 2006 DEM
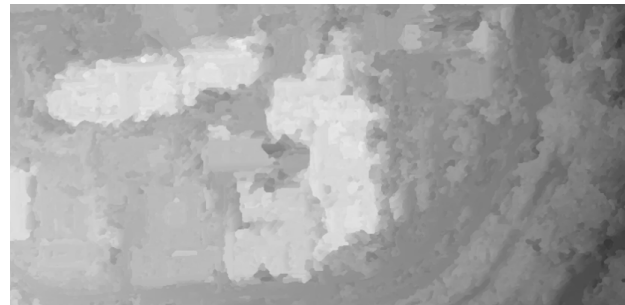
(b) 2010 DEM
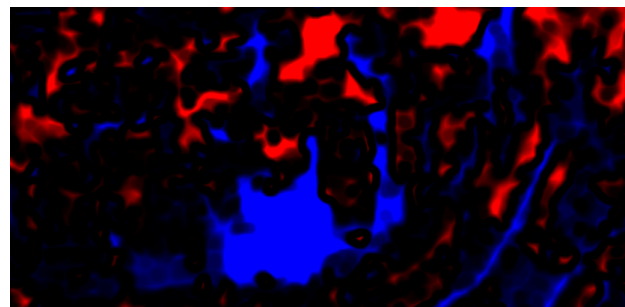
(c) Blurred difference

(d) Thresholding

(e) Erosion

Figure 2: Input DEMs and filtering. Constructions are in blue and destructions in red.

these components). Each component is then analysed using the theory of homological persistence in order to study the compactness of all possible topological subcomponents. The theory of homological persistence is briefly presented in Section 2.2, the definition of a compactness criterion is given in Section 2.3 and the application of both to shape based change detection in Section 2.4. Finally results are given and discussed in Section 3 and conclusions drawn in Section 4.

2 METHOD

2.1 Filtering

We propose to filter the difference map both before and after splitting as shown in Figure 2. A general idea that we push forward in this paper is that the full shape of the change needs to be analysed. By "shape of the change" we mean both the altimetric shape (shape of the difference map over the area of the detected change) and planimetric (shape of the detected change area as seen on a map). Blurring the DEM difference before splitting allows to give more importance to small changes with high altimetric difference. This is performed by convolution of the DEM difference with a Gaussian (Figure 2(c)) which standard deviation should be adapted to the DEM resolution (typically 4 times the pixel size). After this blurring, the DEM difference is split (Figure 2(d)) into positive and negative components. Constructions ($\Delta > 0$) and destructions ($\Delta < 0$) can be separated in this step. Finally, a morphological erosion is performed (Figure 2(e)) to remove thin differences often occurring near building edges. The resulting components should not be interpreted as changes but as areas of potential changes that will be analysed in further steps using homological persistence (Section 2.2) and generalized compactness (Section 2.3).

2.2 Homological Persistence

Homological persistence is a very general study of the evolution of the topology of the level sets of functions as the level varies. We will only present here the special case of 2D functions over the real space ($f : \mathbb{R}^2 \to \mathbb{R}$) and only study the first homology group (of order 0), which represent the connectedness of components. In 2D there is also a second group (of order 1) representing the loops. For the general theory of homological persistence, we point the interested reader to (Cohen-Steiner et al., 2005).

Homological persistence relies on detecting topological events that occur when there is no strict 1-to-1 mapping (by inclusion) between the connected components of:

$$f_a = \{x \in \mathbb{R}^2 | f(x) \le a \in \mathbb{R}\} \quad (1)$$

between two different levels $a_1 < a_2$. In our special case, this may only happen when:

1. A component appears between a_1 and a_2 (minimum of f).

2. Two components merge between a_1 and a_2 (saddle point of f).

For a function f defined over a set S of pixels p, this will be done by constructing a *persistence* tree. Each node and leaf of the tree is a connected components $C_i \subset I$ defined by the set of its pixels and having a unique identifier i. The tree is built by progressively adding all pixels p in descending or ascending order of $f(p)$ and maintaining an identifier map $Id(p)$. Updating the tree
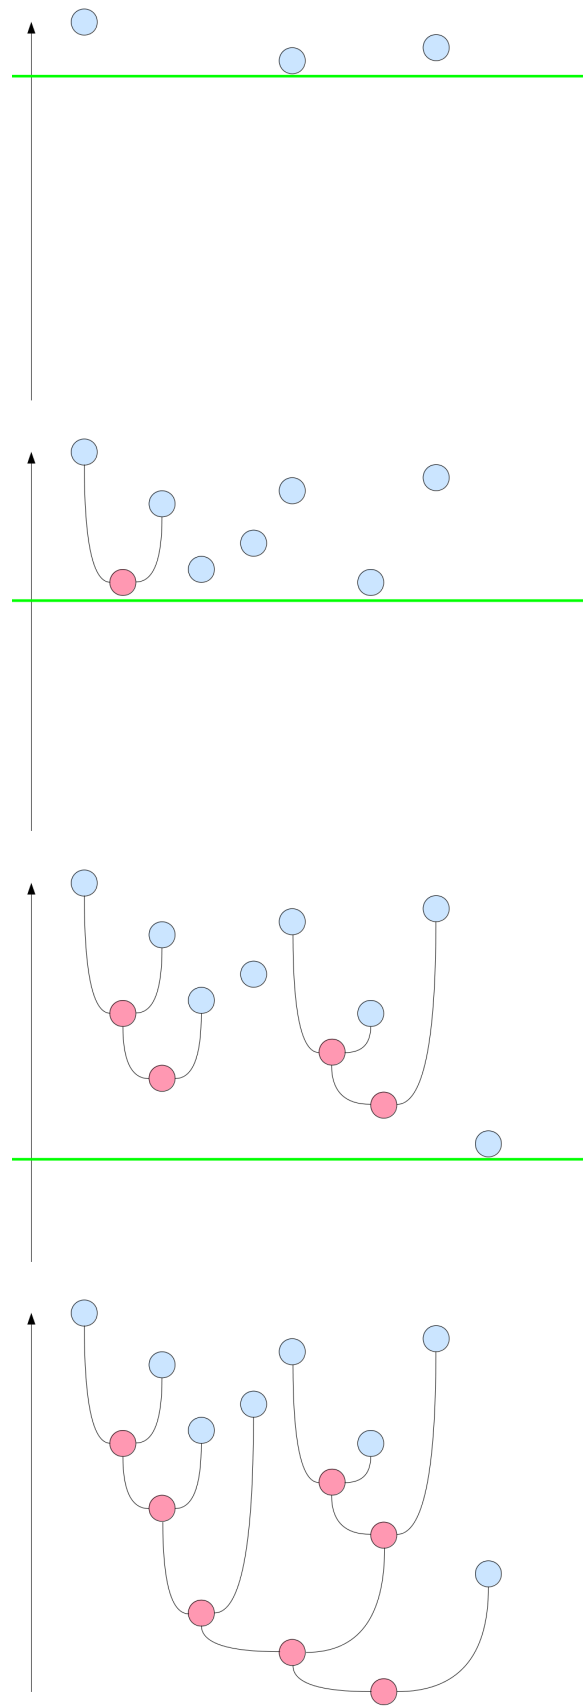

Figure 3: Representation of *persistence* tree computation: red leafs correspond to the apparition of a component, blue nodes to the merge of two components, green line to the current level.

when adding a pixel p depends on the number of different existing active components in its neighbourhood (number of different identifiers in the neighbourhood of p in the identifier map i):

1. 0: creation of a new active component C_i containing only p and setting $Id(p) = i$. Such components are leafs of the persistence tree, they do not have ancestors.

2. 1: growing the neighbouring component C_i by adding p to its set of pixels and setting $Id(p) = i$.

3. $n > 1$: merging the neighbouring components $C_{i_1}, ..., C_{i_n}$ by creating a new component C_j containing all the pixels of all the merged components. Set $Id(p) = j$ and $Id(q) = j \forall q \in C_{i_1} \cup ... \cup C_{i_n}$. Declare $C_{i_1}, ..., C_{i_n}$ as ancestors of C_j. They are now inactive (they won't change any more) as their identifiers have all been replaced by j in the identifier map.

This construction finishes when all the pixels in S have been added. At that point, there is only one active component containing all the pixels of S. The process is illustrated in Figure 3 where the nodes are placed at the creation level of the corresponding component.

The persistence tree represents the topological structure of the level sets of the function, and its construction is parameter-free. When we analyse the DEM difference, the changes that we want intuitively a change detection method to find can be at any level of this tree. Thus we will analyse all the nodes of the tree in order to find the changes of interest, based on several geometric criteria, for instance:

1. Area of the change: number of pixels of the component

2. Maximum, median or average height of the change: maximum, median or average Δz between pixels of the components, or equivalently level difference between the analysed node and the highest leaf over it.

3. Compactness: see Section 2.3.

The first three indicators are very easy to compute (they can be updated each time a pixel is added to a component) and are various indicators of the size (both altimetric and planimetric) of the change. In order to make our geometric analysis finer, we propose, as (Chaabouni-Chouayakh et al., 2010), to use a compactness operator. Compared to this work, our contribution is double:

1. The compactness operator is used to analyse nodes at all the levels of the persistence tree instead of a single arbitrary level.

2. We define a generalized compactness to overcome some limitations of the classical compactness.

2.3 Generalized Compactness

(Chaabouni-Chouayakh et al., 2010) use the classical compactness indicator defined by:

$$\mathcal{C}(C) = \frac{2\sqrt{\pi \mathcal{A}(C)}}{\mathcal{P}(\mathcal{C})} \in [0, 1] \qquad (2)$$

where $\mathcal{A}$ and $\mathcal{P}$ are the area and perimeter of the component C. It is maximum (1) for disks, close to 0 for very complex/elongated

shapes, and 0 for 1D shapes (that have a perimeter but null area). These properties are useful to filter out elongated and complex components, but suffers a strong limitation in our case where components are rasterized (represented by a set of pixels). Indeed rasterization implies that smaller components are more compact than larger ones (it is hard to be complex or elongated is you have few pixels). In other terms, because the pixel size is fixed, compactness is sensitive to scale. To overcome this limitation, we propose the following "generalized" compactness :

$$\mathcal{C}_\alpha(C) = \sqrt{\frac{2(2\pi)^\alpha \mathcal{A}(C)}{\mathcal{P}(\mathcal{C})^{\alpha+1}}} \qquad (3)$$

the constants being chosen such that a disk of radius r has a compactness of $r^{(1-\alpha)/2}$. The $\alpha = 1$ case is the classical compactness criterion which favours small components. Conversely, $\alpha = 0$ favours too much large components so α should be chosen in the $[0, 1]$ interval.

2.4 Application to shape based change detection

We now need an indicator of what is a good geometry for a change, and consider that the answer to this question is very application dependent. In our case, we were mainly interested in building, so we chose a quality estimator of the form:

$$\mathcal{Q}(C) = \mathcal{C}_{0.5}(C)\Delta_z^{av}(C)$$

$\mathcal{C}_{0.5}$ has similar values for small compact components and larger more complex components, which characterizes well the planimetry of buildings, and the second term favours changes for which the altimetric difference is high. This quality criteria combines the quality of the planimetric and altimetric shape of the potential changes.

The selection of components in the persistence tree based on quality criterion $\mathcal{Q}$ is then done in a greedy manner:

1. Compute $\mathcal{Q}$ for each node of the tree

2. While the quality of the best component $\mathcal{Q}(C_{best})$ is above a threshold:

   (a) Add C_{best} to the list of detected changes

   (b) Remove C_{best}, all its ancestors (components containing C_{best}) and successors (components included in C_{best}) from the tree

   (c) Find the new best component in the rest of the tree.

This algorithm allows for a much finer geometric analysis of changes than without using the multilevel approach based on persistence. For instance, the two components of our running example are well separated as shown on Figure 4(a) without requiring a threshold tuning.

Vegetation however still causes problems: because DEMs are rather bad in vegetation, the resulting DEM differences have very random shapes some of which are good according to our quality criteria. Figure 4(b) shows that our algorithm, and especially the generalized compactness criterion allows to remove most false alarms in the vegetation, but some are compact enough to remain. Thus we used the vegetation detection of (Iovan et al., 2008) to remove vegetation found coherently at the two dates. This way, we remove false alarms in what is detected as vegetation at the two dates, but keep changes corresponding to the replacement of a tree with non vegetation (tree cut to build something) or the inverse (something built replaced by trees) which are often meaningful cues.
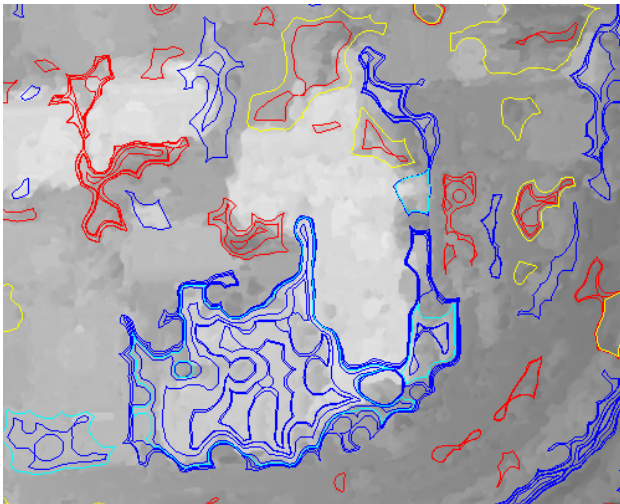
(a) On our running example: the two components circled in Figure 1 are well separated


(b) On the vegetation: most complex components are filtered out. Only the components with compact shapes remain and will need a vegetation mask to be filtered out.

Figure 4: Result of the greedy search: all potential changes are encircled in red (destructions) and blue (constructions), selected changes are in yellow (destructions) and cyan (constructions).

3 RESULTS AND DISCUSSION

A performance study of this algorithm has been conducted on two test areas covering around $2km^2$ near the cities of Phoenix (USA) and Toulouse (France). The "Phoenix" area covers an industrial zone so most of its area is covered with industrial buildings and parking lots with vehicles of various sizes. The "Toulouse" area is in the suburbs and contains a wide variety of typologies: individual and collective housing, large buildings (hospital and school), forest, fields,... The data available on Phoenix was panchromatic satellite imagery: Ikonos in 2008 with 80cm GSD and Worldview2 in 2011 with 46cm GSD. For Toulouse, we used 25cm GSD aerial photographs from 2006 and 70cm GSD Pléiades satellite images from 2010. The four DEMs have been made using the MicMac software (Pierrot-Deseilligny and Paparoditis, 2006), with a standard parameterization.

Results are satisfactory even in difficult cases. Figure 5 shows the correct detection of new buildings, and Figure 6 shows that even (large) vehicles groups can be detected. All the signifi-
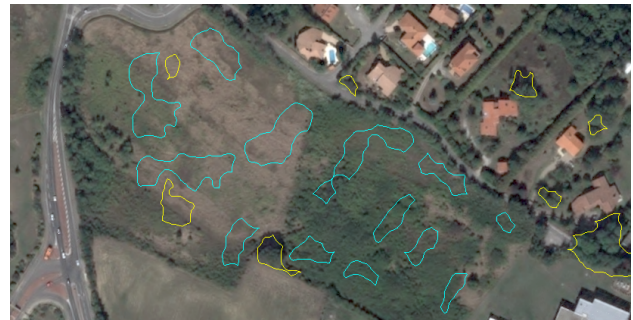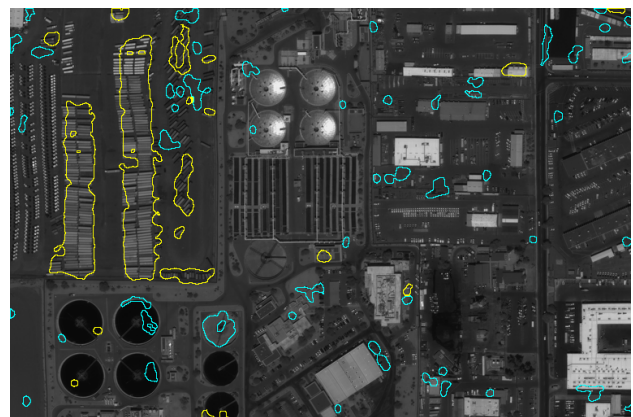

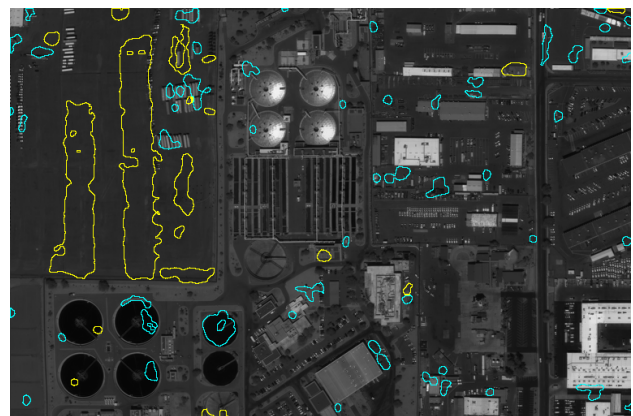(a) 2006 orthophoto with detections encircled

(b) 2010 orthophoto with detections encircled

Figure 5: Constructions correctly detected over the Toulouse area by our method (cyan=constructions, yellow=destructions).

(a) 2008 panchromatic image with detections highlighted

(b) 2011 panchromatic image with detections encircled

Figure 6: Result of the change detection over a part of the Phoenix area (cyan=constructions, yellow=destructions).

Area	Precision	Recall
Phoenix	$160/224 = 71.4\%$	$160/189 = 84.65\%$
Toulouse	$142/231 = 61.5\%$	$142/153 = 92.8\%$

Table 1: Performance of our algorithm on the two selected areas around Toulouse

cant changes are detected, and the compactness criteria eliminates most artefacts. False alarms mainly come from errors in the input DEMs. An interpreter (specialized in change detection from satellite imagery) has validated the results of our algorithm (separating true and false detections) and noted the areas where real changes had not been detected, which allowed us to compute the *Precision* (ratio of good detections over the total number of detections) and *Recall* (ratio of good detections over total number of real changes) which are given on Table 1 for the two areas.

Thanks to the NDVI based vegetation mask, most false alarms in vegetation have been removed. Many false alarms remain in areas where vegetation is not detected consistently between the two dates. It is often hard to know if this is really due to a change between vegetation and non vegetation classes that is not interpreted as meaningful by the operator or to errors in the vegetation masks. Most of the other false alarms come from errors in the DEM computation, and in particular around discontinuities that are often smoothed by the regularization term of the surface reconstruction method. This problem is even increased in our case where the satellites were viewing the scene under a rather high angle (away from NADIR). Things are even worse when the viewing direction is very different at the two dates (which was the case on Phoenix), or when the acquisition typology is different (aerial vs spatial in Toulouse). These considerations are general to all DEM based change detection approaches and not only ours. Finally, it has to be noted that the results were evaluated by professional image interpreters, for which the notion of meaningful changes, in an operational definition, is often quite far from varying heights as we try to define change in this work. Thus many false alarms and missed detections come from this differences of point of view. A tree replaced by bare ground will be detected by our algorithm but often omitted by the interpreter.

Concerning the parametrization of the method, the only parameters to tune are the two parameters of the filtering, the α parameter of generalized compactness and the threshold on quality. The first two are very standard while the last two are very specific to our method and define what the correct shape for a change is. The main contribution of this paper is not the shape analysis that uses these two parameters (even if the generalized compactness is novel and useful) but the (parameter free) persistence tree construction.

4 CONCLUSIONS AND FUTURE WORK

Homologic persistence and the notion of persistence tree allow for a multi-level analysis of the geometry of changes that is not sensitive to the choice of a threshold. Generalized compactness has proved a useful geometric indicator if the change detection focuses mainly on buildings as it allows for the rejection of complex shapes without favouring smaller components too much. The results are highly satisfactory and lead us to think that this algorithm could be used for production purposes, as long as the quality of the DEMs used is sufficient. Enhancing DEM quality is therefore the best lead for improvement. Obviously, the method will also perform better is the DEMs have been acquired in similar conditions (time of day and year, viewing angle, resolution, same sensor, ...) and if the same method (and same parameters for this method) is used.

Future work on the method itself will probably focus on a more global optimization (rather than greedy search) to select the optimal subset of components from the persistence tree, and to integrate other cues than purely geometric ones (using edges extracted from the images and image based classifications at the two dates). Finally, we could exploit the results to use our persistence tree concept along with a supervised learning method instead of simply thresholding a quality criterion.

REFERENCES

Chaabouni-Chouayakh, H., Krauss, T., dAngelo, P. and Reinartz, P., 2010. 3d change detection inside urban areas using different digital surface models. In: Paparoditis N., Pierrot-Deseilligny M., Mallet C., Tournaire O. (Eds), IAPRS, Vol. XXXVIII, Part 3B.

Champion, N., Boldo, D., Pierrot-Deseilligny, M. and Stamon, G., 2011. 2d change detection from satellite imagery: Performance analysis and impact of the spatial resolution of input images. In: Proc. of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Vancouver, Canada.

Cohen-Steiner, D., Edelsbrunner, H. and Harer, J., 2005. Stability of persistence diagrams. In: Proceedings of the twenty-first annual symposium on Computational geometry, SCG '05, ACM, Pisa, Italy, pp. 263–271.

Guérin, C., Binet, R. and Pierrot-Deseilligny, M., 2012. Détection des changements d'élévation d'une scène par imagerie satellite stéréoscopique. In: Reconnaissance des Formes et Intelligence Artificielle, RFIA 2012, Lyon, France.

Hebel, M., Arens, M. and Stilla, U., 2011. Change detection in urban areas by direct comparison of multi-view and multi-temporal als data. In: Proceedings of the 2011 ISPRS conference on Photogrammetric image analysis, PIA'11, Springer-Verlag, Munich, Germany, pp. 185–196.

Iovan, C., Boldo, D. and Cord, M., 2008. Detection, characterization and modeling vegetation in urban areas from high resolution aerial imagery. EEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 1(3), pp. 206–213.

Kuo-hsin, H., Liu, J., Cheng, D. and Yu, M., 2006. Terrain change detection combined photogrammetric dem. Perspective c, pp. 1–6.

Pierrot-Deseilligny, M. and Paparoditis, N., 2006. A multiresolution and optimization-based image matching approach: An application to surface reconstruction from spot5-hrs stereo imagery. In: International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. 36 (Part 1/W41), Ankara, Turkey.

Sasagawa, A., Watanabe, K. N. S., Koido, K., Ohno, H. and Fujimura, H., 2008. Automatic change detection based on pixel-change and dsm-change. In: The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVII. Part B7, Beijin, China, pp. 1645–1650.

Tian, J., Reinartz, P., dAngelo, P. and Ehlers, M., 2013. Region-based automatic building and forest change detection on cartosat-1 stereo imagery. {ISPRS} Journal of Photogrammetry and Remote Sensing 79(0), pp. 226 – 239.

Wheaton, J., Brasington, J., Darby, S. and Sear, D., 2010. Accounting for uncertainty in dems from repeat topographic surveys: improved sediment budgets. Earth Surf. Process. Landforms 35, pp. 136–156.