Ranking evolution maps for Satellite Image Time Series exploration: application to crustal deformation and environmental monitoring

Nicolas Méger, Christophe Rigotti, Catherine Pothier, Tuan Nguyen, Felicity Lodge, Lionel Gueguen, Rémi Andréoli, Marie-Pierre Doin, Mihai Datcu

Nicolas Méger, Christophe Rigotti, Catherine Pothier, Tuan Nguyen, Felicity Lodge, et al.. Ranking evolution maps for Satellite Image Time Series exploration: application to crustal deformation and environmental monitoring. Data Mining and Knowledge Discovery, 2019, 33 (1), pp.131-167. 10.1007/s10618-018-0591-9 . hal-01898015

HAL Id: hal-01898015
https://hal.science/hal-01898015

Ranking Evolution Maps for Satellite Image Time Series Exploration – Application to Crustal Deformation and Environmental Monitoring[⋆]

(Sept. 2018)

Nicolas Méger[1], Christophe Rigotti[2], Catherine Pothier[3], Tuan Nguyen[1], Felicity Lodge[1], Lionel Gueguen[4], Rémi Andreoli[5], Marie-Pierre Doin[6], and Mihai Datcu[7]

[1] Université Savoie Mont Blanc, Polytech Annecy-Chambéry, LISTIC, B.P. 80439, Annecy-le-Vieux, F-74944 Annecy Cedex, France
`nicolas.meger, hoang-viet-tuan.nguyen, jocelyn-felicity-mary@univ-smb.fr`
[2] Univ Lyon, INSA-Lyon, CNRS, INRIA, LIRIS, UMR5205, F-69621, Villeurbanne, France
`christophe.rigotti@insa-lyon.fr`
[3] Univ Lyon, INSA-Lyon, CNRS, LIRIS, UMR5205, F-69621, Villeurbanne, France
`catherine.pothier@insa-lyon.fr`
[4] Uber Technologies, 400 Centennial Parkway, Louisville, CO 80027, USA
`lgueguen@uber.com`
[5] Bluecham SAS, 98807 Nouméa, Nouvelle-Calédonie, France
`remi.andreoli@bluecham.net`
[6] Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, IRD, IFSTTAR, ISTerre, F-38000 Grenoble, France
`marie-pierre.doin@univ-grenoble-alpes.fr`
[7] German Aerospace Center (DLR), Remote Sensing Technology Institute, Oberpfaffenhofen, D-82234 Weßling, Germany
`mihai.datcu@dlr.de`

Abstract. Satellite Image Time Series (SITS) are large datasets containing spatiotemporal information about the surface of the Earth. In order to exploit the potential of such series, SITS analysis techniques have been designed for various applications such as earthquake monitoring, urban expansion assessment or glacier dynamic analysis. In this paper, we present an unsupervised technique for browsing SITS in preliminary explorations, before deciding whether to start deeper and more time consuming analyses. Such methods are lacking in today's analyst toolbox, especially when it comes to stimulating the reuse of the ever growing list

---

[⋆] Funding for this project was provided by a grant from la Région Auvergne-Rhône-Alpes (Tuan Nguyen's grant). It was also supported by the PHOENIX ANR-15-CE23-0012 grant of the French National Agency of Research, and benefited from a Centre National de la Recherche Scientifique (CNRS) "Défi Mastodons" funding. Catherine Pothier and Christophe Rigotti are members of Laboratoire d'Excellence Intelligence des Mondes Urbains (LabEx IMU, ANR-10-LABX-0088) that provided complementary support.

of available SITS. The method presented in this paper builds a summary of a SITS in the form of a set of maps depicting spatiotemporal phenomena. These maps are selected using an entropy-based ranking and a swap randomization technique. The approach is general and can handle either optical or radar SITS. As illustrated on both kinds of SITS, meaningful summaries capturing crustal deformation and environmental phenomena are produced. They can be computed on demand or precomputed once and stored together with the SITS for further usage.

Keywords: Satellite Image Time Series, Summarization, Swap Randomization, Mutual Information, Crustal Deformation, Environmental Monitoring

1 Introduction

A Satellite Image Time Series (SITS) is a set of satellite images acquired at different dates and covering the same geographical area. It contains huge quantities of multidimensional information providing a great potential source of knowledge about the surface of the Earth. Extracting information from these data, which incorporate both spatial and temporal dimensions, is challenging and useful for many applications such as urban expansion assessment (e.g., Marin et al, 2015a; Su et al, 2014; Liu et al, 2012; Cauwels et al, 2014; Duede and Zhorin, 2016), glacier dynamic analysis (e.g., Akbari et al, 2014; Fahnestock et al, 2016; Tedstone et al, 2015), snow cover mapping (e.g., Schellenberger et al, 2012; Crawford et al, 2013), forest mapping (e.g., Zhu and Liu, 2014; Quegan et al, 2000), earthquake monitoring (e.g., Marin et al, 2015b; Wang et al, 2015), coastline detection (e.g., Alonso-Gonzalez et al, 2012; Gonçalves et al, 2014), or soil erosion monitoring and prediction (e.g., Amitrano et al, 2015; Carvalho et al, 2014).

Many SITS have been prepared for such applications and are the result of resource consuming efforts. Most of the time, building a SITS does not mean simply downloading a set of files from an image repository, but requires a design process based on several trade-offs and expert choices. Many questions have to be addressed, such as *“Which bands should be retained among the available ones, or which synthetic band/index should be computed?”, “Taking into account the different spatial resolutions and acquisition dates available, which time frame should be selected?”*. In addition, various other resource consuming tasks must be performed, such as the selection of the images, their co-registration, and also, when necessary, other processes such as the handling of image quality degradation (e.g., atmospheric perturbations).

With the continuous development of acquisition means (e.g., the recent Sentinel satellites of the Copernicus programme) and the extension of open data policies, the number of prepared SITS is likely to have an even more rapid growth. A SITS built to analyze a particular phenomenon is also likely to be useful for different studies in the same domain, or even in other domains. Thus, an important challenge is the development of multi-domain repositories for SITS cross-sharing, so as to ease the use of these series in different studies and to make their preparation efforts beneficial to a wide community. A key aspect of such repositories is their ability to support SITS retrieval, and this, by combining two complementary levels, the first level being the querying of the metadata and the second the browsing of the spatiotemporal information. Fulfilling these needs requires the following questions to be answered. 1) How to find SITS that could be interesting, using selection criteria like location, acquisition band, date, or other various metadata? 2) Then, after having preselected a few SITS, how to quickly scan the main spatiotemporal phenomena they encode, in order to decide whether it is worth downloading the corresponding datasets and starting a more detailed analysis with dedicated techniques? The first question is not simple and has no naive answer, but it can be tackled using current databases, geographical information systems and web service approaches. In contrast, the second question is still a largely open research problem.

We here focus on the second question and present an unsupervised method for summarizing SITS to provide the means to quickly disentangle and browse their spatiotemporal content. To this aim, presenting an overview of a SITS as a video made of the successive images of the series would not be sufficient. Indeed, even if such a video were useful, it would still depict the combined result of the phenomena overlapping in space and time. An important property of the summarization would therefore be its ability to exhibit the main phenomena in isolation.

1.1 SITS analysis approaches

No established method is reported in the literature to build such SITS summaries, but many contributions have been made to analyze SITS and mine spatiotemporal data. Among them, change detection techniques are good candidates when it comes to find variations. They have been proposed to exploit the temporal information with the aim of producing change maps. These maps can be computed efficiently at the pixel level (e.g., Coppin et al, 2004; Lu et al, 2004; Krylov et al, 2013; Rokni et al, 2015), the texture level (e.g., Li and Leung, 2001; Ilsever and Ünsalan, 2012) or the object level (e.g., Bontemps et al, 2008; Lu et al, 2016). They depict changes observed for a scene by comparing images acquired on two dates or more. A change map is likely to be easy to read, but is computed for changes only and prior information about the type of changes is required. For example, one may want to underline abrupt changes due to floods, earthquakes or anthropic activities (e.g., Inglada et al, 2003; Dogan and Perissin, 2014) while others may consider gradual changes such as biomass accumulation (e.g., Vina et al, 2004; Kayastha et al, 2012).

A second family of interesting candidates are the spatiotemporal clustering approaches, since they are designed to analyze and depict a complete SITS. Here, the feature vector associated with an entity to be clustered is the whole vector of values associated with this entity over time, leading to a clustering that must be performed in a high dimensional space. As analyzed by Aggarwal et al (2001), such clusterings can be difficult to interpret and require careful parameter settings (e.g., Gueguen and Datcu, 2007; Gallucio et al, 2008). This curse of dimensionality is commonly tackled by three means: the simple aggregation of dimensions, the definition of distances dedicated to the domain, or the use of more complex application dependent transformations of space. For instance, Petitjean et al (2012) use an adaptation of the Levenshtein distance to measure the similarity between the sequences of pixel value changes, and Nezry et al (1995) propose feature vectors containing aggregated values associated with pixels over time (e.g., average, minimum, maximum). Other methods, such as the one designed by Héas and Datcu (2005), extract and select features (e.g., textural, spatial, spectral) using stochastic models to reduce the dimensionality, and cluster the data in this lower-dimensional space. This clustering is completed by an unsupervised learning procedure to build a graph of dynamic cluster trajectories, resulting in a multi-level hierarchical model.

Another important family of tools for finding regularities over space and time are the data mining techniques based on the extraction of spatiotemporal patterns. These approaches search for patterns that are most of the time extensions of patterns defined along a single dimension, such as the *sequential patterns* proposed by Agrawal and Srikant (1995) or the so called *episodes* defined by Mannila et al (1997). These patterns are indeed good candidates for summarization purposes since they can handle an abstract encoding of the data by using symbols to represent states, objects, locations or intervals of pixel values. Let us mentioned a few kinds of spatiotemporal patterns that have been proposed in this family. For example, patterns representing general frequent sub-trajectories of an object are mined by Cao et al (2005), while Cao et al (2007) extract sub-trajectories that also have to be periodic. In the work of Gudmundsson et al (2007), objects are not considered separately, and a pattern is a group of objects corresponding to some behavior, for instance sharing a common motion (direction, speed) at a given date within the same region of space. Some other approaches do not focus on object displacements, as for instance the one of Alatrista Salas et al (2012) where evolutions of predefined zones (e.g., cities, districts, roads) and their neighborhoods are studied. In the paper of Huang et al (2008), general spatiotemporal relations such as *"event type B occurs closely to event type A in space, but after it over time"* are also extracted according to a notion of spatiotemporal neighborhood that depends on predefined classes of objects. Contrary to most of the aforementioned pattern-based approaches for which no use on SITS has been reported so far, frequent evolution patterns such as *"symbol B occurs after symbol A and is followed by symbol C in sequences of pixel values"* have been defined and exploited to describe SITS at the pixel level in an unsupervised way by Julea et al (2011) or Petitjean et al (2011). These patterns have shown their interest for capturing evolutions. Nevertheless, most of the time, they are too numerous to be quickly browsed by users.

1.2 SITS summarization proposal

In this paper, a method for building summaries of SITS for quick human browsing is presented. It is unsupervised and can be seen as a preliminary step toward the automatic construction of thumbnails reflecting the information contained in a SITS. It attempts to separate the spatiotemporal phenomena and outlines where they occur in space and time. Such summaries are intended to provide an expert with evidence of the presence of phenomena (e.g., subsidence along volcano flanks, seasonal variations of coastal vegetation, development of anthropic activities, modifications of river mouths and lake shores) and help to quickly decide whether or not to start more complex and costly dedicated analyses, depending on the phenomena the expert is interested in.

The summaries are built using an objective entropy-based information measure and the evolutions of the pixel values. The method is neither specific to certain classes of objects nor to certain evolutions (e.g., contour changes, displacement fields). Furthermore it is general enough to be applied to both radar and optical SITS and includes a simple but effective parameter setting procedure. Its application to real data shows that it can handle the common sources of SITS quality degradation (e.g., artifacts, clouds, sensor defects), and irregular time stamps (different time intervals between images). This means that the method does not require a dedicated pre-processing that could introduce bias in the series, such as the application of masks to remove regions covered by clouds (leading to unbalanced representation of regions), or the insertion of intermediate images interpolated from other images in order to fill gaps in the series (introducing auto-correlation bias). The method takes as input a set of evolution patterns as defined by Julea et al (2011). These patterns have two main advantages. Firstly, they have been shown to be effective in capturing meaningful evolutions on both optical and radar SITS, and thus enable to develop a method that is general enough to be applied on both kinds of SITS. And secondly, these patterns incorporate a constraint for focusing on evolutions that are coherent in space. This latter criterion is particularly well suited to geoscience studies. However, if one is not interested in such a coherence, and in handling both optical and radar data, then the method could possibly be adapted/extended to use other evolution patterns such as the ones defined by Petitjean et al (2011).

In the preliminary step of the method, a dedicated spatiotemporal localization map is computed for each evolution pattern. This is followed by the main technical part of our contribution, that is a ranking for finding the most informative core maps which are to be retained to form the summary of the SITS. The overall principle is to compare the maps obtained from the original SITS to the maps obtained from a randomized version of the SITS, and to focus on two opposite kinds of maps: the ones that are observed on both the SITS and the randomized data, and the ones that are strongly affected by the randomization. Here, the key intuition is twofold: 1) on the one hand, a map that is hard to destroy by randomization denotes a strong spatiotemporal structure of the SITS and should be retained, 2) on the other hand, a map that is strongly modified corresponds to a structure that is not present in random data, but that is likely to be specific to the SITS, and thus it should be also selected to be part of the summary. To implement this process, a randomization procedure for SITS extending the Boolean matrix randomization scheme (Cobb and Chen, 2003; Gionis et al, 2007) is proposed and the Normalized Mutual Information measure (Cover and Thomas, 1991) is adapted to compare the map information content. Some preliminary results obtained using these principles were presented by Méger et al (2015).

This paper is structured as follows. After having recalled preliminary definitions in Section 2, the SITS summarization method is presented in Section 3. Section 4 illustrates the application of the method to crustal deformation and environmental monitoring. Section 5 concludes and gives directions for future work.

2 Preliminary definitions

In this section, we recall some preliminary definitions and the concept of *Grouped Frequent Sequential patterns (GFS-patterns)* defined in Julea et al (2011) as an extension of the sequential patterns proposed by Agrawal and Srikant (1995) to the domain of SITS.

2.1 Symbolic SITS

A SITS is built by acquiring n satellite images having the same size at different dates (no specific temporal sampling is assumed) over the same area. It can be either an optical or a radar SITS. More precisely, a SITS can originate from two different kinds of satellite sensors. Either sensors are passive, i.e., they generally measure the sunlight radiation reflected from the Earth in the visible, infrared, thermal infrared and microwave parts of the electromagnetic spectrum, or they are active, i.e., they measure radar or laser signals emitted by the sensor itself and reflected from the Earth. Data originating from passive sensors are often referred to as optical data while data generated by radar and, thus active sensors, are termed radar data. Radar and optical SITS can be used to form SITS expressing synthetic bands such as displacement in the line of sight or vegetation indices.

A *symbolic SITS* is a representation of a SITS in which pixel values are encoded by symbols of a discrete domain $\mathcal{D}$. This encoding can be obtained by equal interval bucketing, equal frequency bucketing (percentiles) or by more sophisticated image quantizations.

In a symbolic SITS, a *pixel evolution sequence* is then a pair $((x, y), \langle \alpha_1, \alpha_2, \ldots, \alpha_n \rangle)$ where (x, y) gives the location of the pixel in space (column and row numbers) and the tuple $\langle \alpha_1, \alpha_2, \ldots, \alpha_n \rangle$ contains the symbols representing the values of this pixel over time, ordered from the first to the last image of the SITS. Missing values are encoded by an extra symbol that does not belong to $\mathcal{D}$, and that will be ignored in the following steps of the method.

A toy symbolic SITS is given in Example 1. It describes the evolution of four pixels through five images using a domain of four symbols $\mathcal{D} = \{1, 2, 3, 4\}$.

Example 1. Symbolic SITS containing four pixel evolution sequences through five acquisitions.

$$\begin{gathered}((0,0),\langle 1,1,4,3,2\rangle),\\ ((0,1),\langle 2,1,4,2,3\rangle),\\ ((1,0),\langle 4,1,3,1,3\rangle),\\ ((1,1),\langle 4,1,1,1,3\rangle)\end{gathered}$$

Note that the symbols are not required to be semantically connected. However the human interpretation can be eased if their choice reflects a natural ordering (e.g., 1 for low pixel values and 4 denoting high values), as done for the experiments reported in Section 4.

2.2 Grouped frequent sequential patterns

Patterns and occurrences An *evolution pattern* β over $\mathcal{D}$ is of the form $\beta = \beta_1 \to \beta_2 \to \ldots \to \beta_k$ with $\beta_1, \ldots, \beta_k \in \mathcal{D}$. Intuitively, a pixel satisfies pattern β if, somewhere in its evolution sequence we found symbol β_1, then sometime later β_2 and so on, before finally observing symbol β_k. More formally, the pattern β occurs in an evolution sequence $((x, y), \langle \alpha_1, \alpha_2, \ldots, \alpha_n\rangle)$ if there exist integers $1 \leq i_1 < i_2 < ... < i_k \leq n$ such that $\beta_1 = \alpha_{i_1}, \beta_2 = \alpha_{i_2}, ..., \beta_k = \alpha_{i_k}$. The integers $i_1, \ldots, i_k$ identify a so called *occurrence* of β. No temporal constraint is set and the elements $\beta_1, \ldots, \beta_k$ of β do not need to occur in strictly consecutive images: the occurrences of these symbols can be separated by one or more images. In addition, different occurrences of a pattern β are not required to be synchronized, i.e., two occurrences of an evolution pattern do not need to be built using symbols located in the same images of the series. For instance, in Example 1, the pattern $1 \to 1 \to 3$ occurs in the evolution of the pixel $(0, 0)$ starting in the first image and ending in the fourth, and it also appears in the evolution of pixel $(1, 0)$ spreading from the second to the last image. This pattern also occurs in the evolution sequence of the pixel located at $(1, 1)$, but for this pixel it has three occurrences since there are three ways to form the part $1 \to 1$ of the pattern for this pixel. Furthermore, a symbol occurrence can belong to several occurrences of the same pattern or of different patterns. For example, when considering pixel $(1, 0)$, the occurrence of symbol 4 can form two different occurrences of pattern $4 \to 1$ and it can also be part of the occurrences of other patterns such as $4 \to 3$.

Support and connectivity A pixel is said to be *covered* by a pattern β if its evolution contains at least one occurrence of β. The number of pixels covered by β is called the *support* of β and is denoted $support(\beta)$. Multiple occurrences of β in a same pixel evolution sequence therefore increment $support(\beta)$ by one only, and the support is then proportional to the area in the SITS where evolution pattern β can be observed.

Simply selecting patterns that cover a sufficient area is not stringent enough for SITS, because the measures related to a phenomenon impacting the surface of the Earth are spatially autocorrelated (Griffith and Chun, 2016). An evolution pattern reflecting such a phenomenon is thus more likely to occur in locations that tend to be adjacent in space. This can be expressed by a local connectivity constraint, imposing that an evolution location is connected on average to a sufficient number of pixels sharing the same evolution. Checking this connectivity requires setting a spatial analysis window. The proposed method relies on the smallest window surrounding a covered pixel, i.e., the window 3×3 centered on the pixel and including its 8 nearest neighbors. If a larger windows were used (e.g., 5×5 neighborhood), it would not account only for pixels adjacent to the center of the window, and thus would lead to a weaker notion of connectivity.

Let p be a pixel covered by β, the local connectivity for β at p is the number pixels that are also covered by β among the 8 nearest neighbors of p. The average of this value over all pixels covered by β is the *average connectivity of* β and is denoted $AC(\beta)$. A high average connectivity thus reflects the tendency of a pattern to cover pixels that are grouped in space.

In Example 1, the pattern $1 \to 1 \to 3$ has five occurrences, but appears in the evolution sequences of only three different pixels, thus its support is equal to 3. Its average connectivity is $AC(1 \to 1 \to 3) = (2+2+2)/3 = 6/3 = 2$.

The *Grouped Frequent Sequential patterns* (GFS-pattern), as defined in Julea et al (2011), are all possible evolution patterns β satisfying a minimun support threshold σ and a minimun average connectivity threshold κ, i.e., such that $support(\beta) \geq \sigma$ and $AC(\beta) \geq \kappa$. In the rest of this paper, they are simply called *patterns* when clear from the context. Notice that these patterns can express various types of evolutions such as abrupt variations (e.g., $3 \to 3 \to 1 \to 1$), gradual changes (e.g, $3 \to 3 \to 2 \to 2 \to 1 \to 1$), constant behaviors (e.g., $1 \to 1 \to 1 \to 1$) or repeated variations (e.g., $1 \to 2 \to 3 \to 1 \to 1 \to 2 \to 3$).

3 SITS summarization: CE-maps and NMI-based rankings of CE-maps

This section presents the method introduced in this paper to summarize SITS in an unsupervised way. The summary is obtained in the form of a small set of representative maps, depicting in isolation evolutions occurring in the SITS. The main technical part of the method is the ranking of the maps to filter the most informative ones with respect to an entropy criterion. This ranking is performed using a randomization procedure and a scoring based on a Normalized Mutual Information measure.

3.1 Core Evolution maps (CE-maps)

For each evolution pattern $\beta = \beta_1 \to \beta_2 \to \ldots \to \beta_k$, a map is built. This map is an image having the same size as the images of the SITS, where all pixels are set to a value of 0 except those covered by the pattern. For each pixel p covered by β, let (x, y) be the location of p and $i_1, \ldots, i_k$ the temporal position of the first occurrence of β found in the evolution sequence of p. Then, the value of the pixel at location (x, y) in the map is set to i_k, i.e., the index of the image in the SITS where β_k occurs. This map simply gives the spatial locations of the occurrences and their respective ending locations over time, while the temporal evolution shape is given by pattern β itself. A color scale is used to represent the strictly positive values of the pixels of the map (earliest ending dates of pattern β), and pixels having a value of 0 (no occurrence of the pattern) are depicted in black. Other temporal indices such as starting indices (i.e., i_1), central indices (e.g., the average of the image indices supporting the pattern occurrence) or occurrence durations (i.e., $i_k - i_1$) are also reasonable alternatives. Starting/ending indices are of special interest as, once visualized, they show the temporal gradients of the beginning/ending of the phenomena. Ending indices have been retained since they clearly depict the dates for which the whole evolutions have been fully observed.

Typical maps are given in Figures 1a and 1b. As detailed in Section 4, these maps were obtained when building the summary of an *Interferometric Synthetic Aperture Radar* Envisat SITS covering Mount Etna between 2003 and 2010. They express Earth surface deformations caused by the volcanic and seismic activities affecting the region. The colors used to build the maps range from red and orange (early dates in the series) to magenta and rose (late dates in the series) according to the color scale given in Figure 1e. As shown in Figure 1c and Figure 1d, the map interpretation can be eased by overlaying them on the Digital Elevation Model (DEM) of the zone.

Simply by observing these maps, and without relying on the precise meaning of the symbols forming the patterns, it can be noticed, for instance, that when evolution pattern $1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 3$ of Figure 1a occurs then it always ends in the same image of the SITS (only one color is present). On the contrary, $1 \to 1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$, depicted in Figure 1b, ends at different time stamps (several colors). More precisely, the map of Figure 1a reflects an evolution ending in the last image of the series (color magenta), whilst the map of Figure 1b shows an evolution ending by the middle of the series (color light blue) and that propagates in space over time till the end of the series (color magenta). Surprisingly, even such simple patterns are sufficient to reflect and isolate different phenomena over space and time. As explained in Section 4, the two maps of Figures 1a and 1b match geologic structures (a lava plateau and fault systems reported in the literature) impacted by a regular global uplift (symbol 1 in the patterns) counterbalanced by temporary deflation phases (symbols 2 and 3).

For the sake of readability of the summary, the disentanglement of the phenomena is then accentuated. To this aim, the method searches for the maps that focus on the most specific areas and evolutions, while still reflecting the shapes of all the evolutions that have been encountered. This is achieved by a maximality-based filtering developed in pattern mining (e.g., Burdick et al, 2005; Gouda and Zaki, 2001; Luo and Chung, 2005) and defined as follows. Let $\beta = \beta_1 \to \beta_2 \to \ldots \to \beta_k$ and $\beta' = \beta'_1 \to \beta'_2 \to \ldots \to \beta'_{k'}$ be two evolution patterns such that $k' < k$, then β' is a *subpattern* of β if there exist integers $1 \leq i_1 < i_2 < ... < i_k \leq k$ such that $\beta'_1 = \beta_{i_1}, \beta'_2 = \beta_{i_2}, ..., \beta'_{k'} = \beta_{i_k}$. Informally, β' is a subpattern of β if it can be obtained from β by removing some symbols.

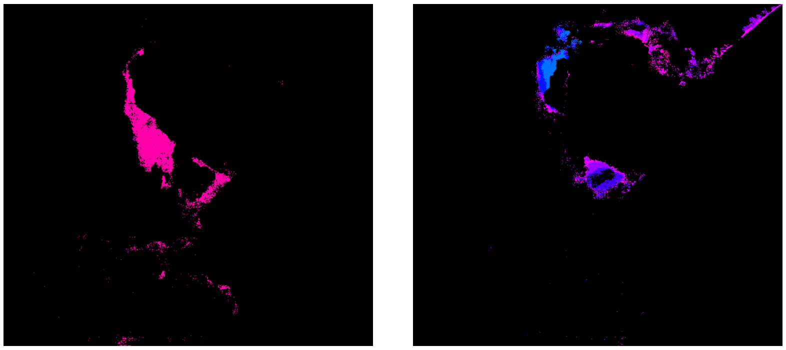

(a) Map of pattern $1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 3$

(b) Map of pattern $1 \to 1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$

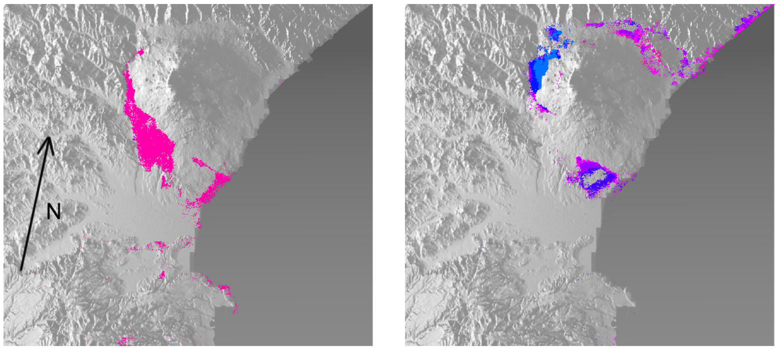

(c) Map of pattern $1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 3$ overlaid with the Mount Etna DEM

(d) Map of pattern $1 \to 1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$ overlaid with the Mount Etna DEM

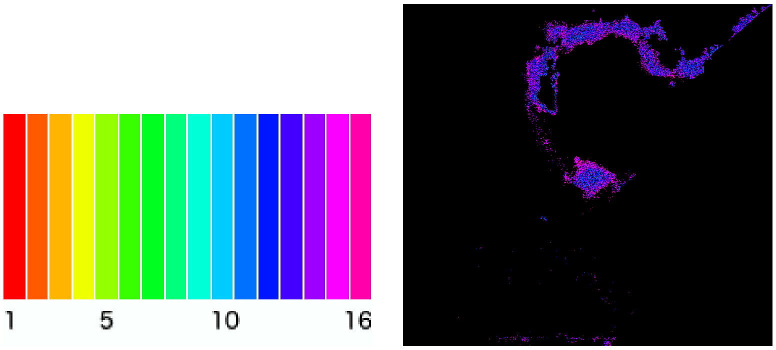

(e) Temporal color scale: from red and orange on the left (occurrences ending early in the series) to magenta rose on the right (occurrences ending late in the series). Numbers denote image numbers, from 1 to 16

(f) Map of pattern $1 \to 1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$ over a randomized SITS

Fig. 1 Maps of two patterns extracted from an InSAR Envisat SITS covering Mount Etna

The *maximal* evolution patterns of a collection of patterns $\mathcal{C}$ are the elements in $\mathcal{C}$ that are not subpatterns of any other pattern in $\mathcal{C}$.

For instance, let $\mathcal{C} = \{1 \rightarrow 3 \rightarrow 2, 1 \rightarrow 3 \rightarrow 1 \rightarrow 2, 3 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 2 \rightarrow 1, 1 \rightarrow 2 \rightarrow 1\}$. In this collection, the first pattern is a subpattern of the second and of the third, the fourth is a subpattern of the third one, and thus the maximal patterns are $1 \rightarrow 3 \rightarrow 1 \rightarrow 2$ and $3 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 2 \rightarrow 1$. Note that maximality is not the same as the notion of *longest pattern*. In this example, $3 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 2 \rightarrow 1$ is the longest pattern, while $1 \rightarrow 3 \rightarrow 1 \rightarrow 2$ is not the longest, but is a maximal one.

The maximal evolution patterns are then used to define the notion of *Core Evolution maps* (CE-maps).

Definition 1. *For a given collection $\mathcal{C}$ of patterns, a CE-map is a map corresponding to a maximal pattern in $\mathcal{C}$.*

When building a summary, selecting maps that are CE-maps helps to direct the attention to phenomena in isolation. This results from two mechanisms. Firstly, a focus on certain evolutions comes directly from the definition of CE-maps as maps associated with the most specific evolution patterns. Secondly, the CE-maps also exhibit specific core spatial locations where the phenomena occur. This behavior with respect to the spatial component can be shown as follows. Let $\mathcal{A}$ be the set of non-black pixels in a CE-map associated to a pattern β. Then, $\mathcal{A}$ is included in the set of non-black pixels of any map of a subpattern of β, because a pixel covered by β is also covered by any subpattern of β. So, $\mathcal{A}$ represents the most restricted area covered by β and its subpatterns.

It is worth noting that the CE-map selection not only helps in the disentanglement of the different phenomena, but also preserves some information about the other maps. Indeed, even if a map corresponding to a highlighted set of pixels $\mathcal{A}'$ and a pattern β' is not retained as a CE-map, then we know that β' is still encompassed by at least one pattern β associated with an existing CE-map. In addition, this CE-map sketches roughly the spatial location of $\mathcal{A}'$ by pointing-out the shared locations among all subpatterns of β. Hence, while focusing on some evolutions and areas, the CE-maps still outline the contents of the other maps, and thus are appropriate for summarization purpose.

3.2 Ranking of CE-maps based on Normalized Mutual Information (NMI)

While the CE-maps are easy to read, as shown in Figure 1, they are too numerous for quick human browsing (hundreds of CE-maps can be obtained, as reported in Section 4). So, an important step is then to focus on a reduced number of CE-maps. This problem is not the same as choosing the CE-maps that could be surprising with respect to some symbol distributions or some other information already known by the user, because, to be representative, the summary should also encompass known and obvious aspects of the SITS content.

In order to remain as unsupervised as possible, the method does not rely on any external indices or models, but is solely based on a scoring related to the low level SITS and CE-map content.

Let $\mathcal{S}$ be a SITS and $\mathcal{M}$ be a CE-map on $\mathcal{S}$ for a pattern β, then the scoring of $\mathcal{M}$ is performed as follows. It consists in comparing $\mathcal{M}$ to the CE-map $\mathcal{M}'$ of β on a SITS $\mathcal{S}'$, where $\mathcal{S}'$ is obtained from $\mathcal{S}$ by breaking both the regions contained within each image and the ordering of the symbols within each sequence. This comparison is used to identify two situations:

- $\mathcal{M}$ and $\mathcal{M}'$ are very different, suggesting that $\mathcal{M}$ denotes phenomena that are not likely to be observed on variant of $\mathcal{S}$, and that the information convey by $\mathcal{M}$ are rather specific to $\mathcal{S}$.
- $\mathcal{M}$ and $\mathcal{M}'$ are very similar, suggesting that $\mathcal{M}$ depicts phenomena that are difficult to hide by modifying $\mathcal{S}$ and that are in some sense strongly represented in $\mathcal{S}$.

In both cases, such a CE-map $\mathcal{M}$ deserves to be used to build a representative summary of $\mathcal{S}$. Thus, the method associates a score with each CE-map $\mathcal{M}$ based on a measure of similarity between the CE-maps $\mathcal{M}$ and $\mathcal{M}'$. Then, the summary is built by selecting a few CE-maps having the lowest scores (low similarities, i.e., CE-maps changing the most from $\mathcal{S}$ to $\mathcal{S}'$), and by also selecting a few CE-maps having the highest scores (high similarities, i.e., CE-maps undergoing the smallest changes).

The two following subsections present how to obtain $\mathcal{S}'$ from $\mathcal{S}$ by mean of a swap randomization technique, and then how to assess the similarity of $\mathcal{M}$ and $\mathcal{M}'$ using a normalized mutual information measure in order to rank $\mathcal{M}$.

Swap randomization of SITS

Changing the symbol connectivity within the i-th image of a SITS $\mathcal{S}$ can be performed by computing a random permutation of the symbols in this image. Unfortunately, this simple solution is too naive and will modify the content of most pixel evolution sequences $s = ((x, y), \langle \alpha_1, \ldots, \alpha_i, \ldots, \alpha_n \rangle)$ at index i. Thus, this will not only change the order within each sequence, but after having replaced each image by its random permutation, all pixel evolution sequences are likely to have changed deeply. Then, except for some specific SITS with particular symbol distributions, all CE-maps will be very different for the new SITS $\mathcal{S}'$.

The problem of breaking the symbol connectivity within each image and the ordering within each sequence can be solved by adapting an operator used for swap randomization of Boolean matrices by Cobb and Chen (2003) and Gionis et al (2007). This operator is designed for randomization purposes, and performs an elementary modification of a 0/1 matrix while preserving the number of 1's in each row and in each column. The key hint is to exchange the values of two cells, while at the same time exchanging the values of two other cells chosen in order to preserve the number of 1's in each row and in each column. The swap randomization of a Boolean matrix is carried out by a series of swap attempts based on this operator.

In the case of a SITS, an elementary swap can be made by exchanging two symbols at locations (x, y) and (x', y') within an image, and at the same time exchanging symbols at the same locations in another image carefully chosen to maintain the frequencies of the symbols in each sequence, and thus breaking only the spatial connectivity and the temporal ordering. More precisely, for a SITS we define an *elementary swap attempt* as follows.

Let i, j be two image indices in a SITS $\mathcal{S}$ such that $i \leq j$, and let $s = ((x, y), \langle \alpha_1, \ldots, \alpha_i, \ldots, \alpha_j, \ldots, \alpha_n \rangle)$ and $s' = ((x', y'), \langle \alpha'_1, \ldots, \alpha'_i, \ldots, \alpha'_j, \ldots, \alpha'_n \rangle)$ be two pixel evolution sequences in $\mathcal{S}$, where i, j, s and s' are chosen randomly such that $\alpha_i = \alpha'_j$ (i.e., the two symbols are the same). Then:

- if $\alpha'_i = \alpha_j$ and $\alpha_i \neq \alpha'_i$, exchange α_i with α'_i and, at the same time, exchange α'_j with α_j.
- if $\alpha'_i \neq \alpha_j$ or $\alpha_i = \alpha'_i$, do nothing (i.e., no real swap made).

So, after a swap attempt, either $\mathcal{S}$ is unchanged, or it contains the modified sequences $s = ((x, y), \langle \alpha_1, \ldots, \alpha'_i, \ldots, \alpha'_j, \ldots, \alpha_n \rangle)$ and $s' = ((x', y'), \langle \alpha'_1, \ldots, \alpha_i, \ldots, \alpha_j, \ldots, \alpha'_n \rangle)$. In this case, since $\alpha_i = \alpha'_j$ and $\alpha'_i = \alpha_j$, then the frequencies of the symbols in each sequence are still the same, and also stay the same in each image, while both the initial ordering of the symbols over time and their location in space have changed.

The complete process starts from an initial SITS S and performs iteratively a sequence of elementary swap attempts, until a sufficient mixing of the data is obtained. With regard to the appropriate number of swap attempts, it is worth noting that no conclusive theoretical result is available in the literature. Intuitively, to be randomized, a large dataset is likely to require more swaps than a smaller one. It is empirically estimated (Gionis et al, 2007) that the number of random swap attempts should be of the order of the size of the dataset. To adopt a conservative setting, the method is applied in Section 4 with the number of swap attempts set to about 20 times the size of the dataset (i.e., 20 × number of pixels per image × number of images).

Note that, as for the swap randomization procedure of Boolean matrices established by Cobb and Chen (2003) and Gionis et al (2007): 1) all swap attempt locations have the same probability of being chosen and can be chosen more than once; and 2) an effective swap can be later undone by another effective swap. As detailed by Gionis et al (2007), the swap attempts that do not result in a change of the dataset (no real swaps made) are necessary to obtain each reachable dataset with the same probability. Indeed, the swap randomized dataset generated from an original dataset can be seen as the result of a random walk on a connected and undirected transition graph. It can be formalized as a Markov chain, where each transition is a swap and each state is a swap randomized dataset. In our case, the vertices of the transition graph are the original dataset and all the datasets that can be reached by swap randomization. The edges are the elementary swap attempts, and the attempts that do not change the data are denoted by self-loops over the vertices. The reversibility of the swaps and the presence of the self-loops result in a uniform stationary distribution for the Markov chain (Cobb and Chen, 2003; Gionis et al, 2007). This ensures that each reachable swap randomized dataset has the same probability to be obtained.

Figure 1f shows the CE-map of the same pattern as the one of the CE-map depicted in Figure 1b, but obtained after applying the randomization procedure. As will be reported in Section 4, this map is among the ones that change the most when compared to their versions on the original dataset. Notice that the occurrences of the pattern have been modified over time as denoted by the changes in the colors, and have also been modified over space resulting in locations where occurrences appear or disappear (black pixels disappearing/appearing).

NMI-based ranking of CE-maps using the randomized SITS

To rank the CE-maps, the next stage is then to assess each CE-map by comparing it to its version on the randomized dataset. More precisely, a CE-map $\mathcal{M}$ of a pattern β on the original SITS $\mathcal{S}$ is compared to the CE-map $\mathcal{M}'$ of β on the SITS $\mathcal{S}'$ obtained by swap randomization of $\mathcal{S}$. At first glance, a CE-map is simply a $n_x \times n_y$ matrix containing ending indices of pattern occurrences and can be represented as an *object* in a $n_x \times n_y$ dimensional space. Since two CE-maps over series $\mathcal{S}$ and $\mathcal{S}'$ have the same size, a simple approach to compare them would be to use a p-norm distance (e.g., Euclidian, Manhattan). In the case of CE-maps, the value of $n_x \times n_y$ corresponds to the size of the images in the SITS and is likely to be large, resulting in a very high dimensional space in which p-norm distances are known to be poor dissimilarity measures as explained by Aggarwal et al (2001). In this situation, a standard method consists in reducing the number of dimensions by using a limited number of derived features to describe the objects. The choice of features is mostly guided by the application, since features capture the aspects of the objects to be used as salient elements in the comparison. It is unclear how to choose these features for CE-maps while preserving the unsupervised aspect of the SITS summarization method, and not introducing additional bias in favor or against some kinds of phenomena in the SITS. For the design of the method, the guideline is here to avoid such transformations to a new feature space if it is not necessary.

Another approach could be to compare two CE-maps by counting the number of locations where they differ, in a way similar to the comparison of strings using the Hamming distance. Such a dissimilarity measure is computed by comparing in turn each location in isolation, and thus is less likely to be biased towards some patterns in the CE-maps. Its drawback for a ranking in a summarization method is that a Hamming-like distance does not capture the information loss due to the missing elements. For instance, if a number z of elements of $\mathcal{M}$ are missing in $\mathcal{M}'$, they result in a dissimilarity value equal to z, regardless of the proportion of information these elements convey in $\mathcal{M}$.

More complex measures such as Jaccard (Jaccard, 1902) or Tanimoto (Tanimoto, 1958) distances have been used to compare objects described by vectors of Boolean attributes. Nevertheless, using these measures would require binarized CE-maps, e.g., by setting black pixels to 0 and other pixels to 1. The resulting comparison would therefore focus on the spatial information and ignore the temporal information. Another measure, considering information contents and termed *mutual information* (Cover and Thomas, 1991), has proven to be appropriate for images and is used by many methods in medical image processing (e.g., Collignon et al, 1995), general image alignment (e.g., Viola and Wells, 1995) and remote sensing imagery (e.g., Chen et al, 2003; Inglada and Giros, 2004; Suri and Reinartz, 2010). One of its main advantages is that, when used on two images, this measure is based solely on the marginal and joint entropies of the images and does not rely on any particular pattern occurrence or any structure in the images. This makes the mutual information very well suited for the comparison of the CE-maps.

The key intuition used to apply mutual information to CE-maps is the same as for images in general. It consists in interpreting the values of the pixels in a CE-map $\mathcal{M}$ as the realizations of a random variable X over a sample space Ω with, in the case of a CE-map, $\Omega = \{0, 1, 2, \ldots, n\}$, i.e., the possible image indices in the SITS or value 0. For two CE-maps $\mathcal{M}$ and $\mathcal{M}'$, interpreted respectively as the realizations of variables X and X', the mutual information of these variables is defined as:

$$I(X, X') = H(X) + H(X') - H(X, X')$$

where $H(X)$ and $H(X')$ are the marginal entropies of variables X and X', and $H(X, X')$ is the joint entropy of the two variables. The mutual information is symmetric (Cover and Thomas, 1991), i.e., $I(X, X') = I(X', X)$, and its informal meaning is simply a measure of the amount of information shared by the two variables.

The entropies are computed as follows, as for instance done by Chen et al (2003). Let $P(X = \omega)$ and $P(X' = \omega)$ be the probabilities that variables X and X' take the value ω, and let $P(X = \omega_1, X' = \omega_2)$ be the joint probability of having $X = \omega_1$ and $X' = \omega_2$, then the entropies are obtained from:

$$\begin{aligned}
H(X) &= -\sum_{\omega \in \Omega} P(X = \omega) \log P(X = \omega) \\
H(X') &= -\sum_{\omega \in \Omega} P(X' = \omega) \log P(X' = \omega) \\
H(X, X') &= -\sum_{\omega_1, \omega_2 \in \Omega^2} P(X = \omega_1, X' = \omega_2) \log P(X = \omega_1, X' = \omega_2)
\end{aligned}$$

using usual conventions for entropy ($0 \log 0$ is assigned the value of 0, and log is the logarithm to base 2).

The probabilities are themselves estimated using the two images, in our case the CE-maps $\mathcal{M}$ and $\mathcal{M}'$. The estimate taken for the marginal probability $P(X = \omega)$ (resp. $P(X' = \omega)$) is the proportion of pixels in $\mathcal{M}$ (resp. $\mathcal{M}'$) having value ω. And for the joint probability $P(X = \omega_1, X' = \omega_2)$ the estimate used is the proportion of the pixels in $\mathcal{M}$ such that the pixel in $\mathcal{M}$ has value ω_1 and the pixel at the same location in $\mathcal{M}'$ has value ω_2.

The mutual information is then interpreted both as a measure of *dependence* between two images and as the amount of information each image contains about the other one (e.g., Pluim et al, 2003). In the summarization method, it is used as a measure of similarity of two CE-maps, based on their information content. To obtain a similarity reflecting the *proportion* of information that is shared by these two CE-maps, an additional normalization step is performed. There exist several definitions of *Normalized Mutual Information* (NMI), see for instance Pluim et al (2003) or Konings et al (2015), and the one retained here is the direct proportion computation:

$$NMI(X, X') = \frac{I(X, X')}{min(H(X), H(X'))}$$

It is based on the upper bound of mutual information given by the property $I(X, X') \leq min(H(X), H(X'))$ that is given by Cover and Thomas (1991) and leads to a NMI ranging from 0 to 1. Notice also that this NMI value does not depend on the base of the logarithm used to compute the entropies.

Finally, the NMI must be restricted to account for the part of the data making sense in the comparison. In a CE-map $\mathcal{M}$, a missing value for the date of occurrence of the pattern is represented by a 0 at the corresponding location in space. These pixels, depicted in black, can be numerous, as for instance in typical CE-maps in Figures 1a and 1b. If such a location still contains no occurrence of the pattern in $\mathcal{M}'$ after randomization of the SITS then it is not considered to be meaningful in the comparison. Thus, to compare $\mathcal{M}$ and $\mathcal{M}'$ using $NMI(X, X')$, the locations where the value is 0 in both CE-maps are discarded, and are not retained as realizations of random variables X and X' when computing the marginal and joint probabilities.

The whole ranking is simply performed as follows. First the swap randomization process is applied to the original SITS $\mathcal{S}$ to obtained $\mathcal{S}'$. Next, for each CE-map $\mathcal{M}$ on $\mathcal{S}$, the corresponding map $\mathcal{M}'$ for the same pattern is computed for $\mathcal{S}'$. The score associated with $\mathcal{M}$ is then the NMI measuring the similarity of $\mathcal{M}$ and $\mathcal{M}'$. Finally, the CE-maps on $\mathcal{S}$ are sorted according to this score.

The CE-maps having a NMI value among the lowest or among the highest ones can then be presented as a summary of the spatiotemporal phenomena that have been isolated in the SITS. More CE-maps can be stored and shown (according to their ranking) to users looking for a more detailed summary.

4 Application to crustal deformation and environmental monitoring

In this section we report the application of the summarization method to crustal deformation and environmental monitoring using two very different datasets: a radar SITS over a volcano, and an optical SITS over a cloudy coastal area. Section 4.1 presents the datasets and their preprocessing, Section 4.2 focuses on parameter settings and Section 4.3 details the results that have been obtained.

4.1 Datasets and preprocessing

The first dataset is a SITS obtained using a multitemporal Interferometric Synthetic Aperture Radar (InSAR) technique on data from the Envisat mission, covering the Geohazard Supersite of the Mount Etna volcano. This SITS is a series of images (598x553) of total phase delays extracted from Envisat ascending tracks (satellite looking eastward). It contains 16 images between early 2003 and summer 2010, with a resolution of about 160 m. For each year of the period, there are two images, one in winter and one in summer. These total phase delay images were produced by temporal inversion of a large series of unwrapped interferograms created using the NSBAS processing chain designed by Doin et al (2011, 2015), including atmospheric correction as also performed by Doin et al (2009).

In these images, positive values correspond to motions away from the satellite along the Line Of Sight (LOS), with respect to the beginning of the series, and negative values correspond to motions towards the satellite. It is important to note that because of the geometric configuration, these displacements contain a part of both vertical and lateral components of the ground motion, i.e. the crustal deformation. Figure 2a shows such an image, where light greys denote positive values and dark greys represent negative values. The satellite ascending track is located outside on the left of the image (parallel to its border). It can be noticed from this image that data can be missing on large areas (depicted here as a uniform grey background), this problem being inherent to the InSAR data and due to the removal of areas subject to lack of coherence. The whole SITS is summarized in Figure 2b by an average LOS velocity map computed using the method of Doin et al (2011).

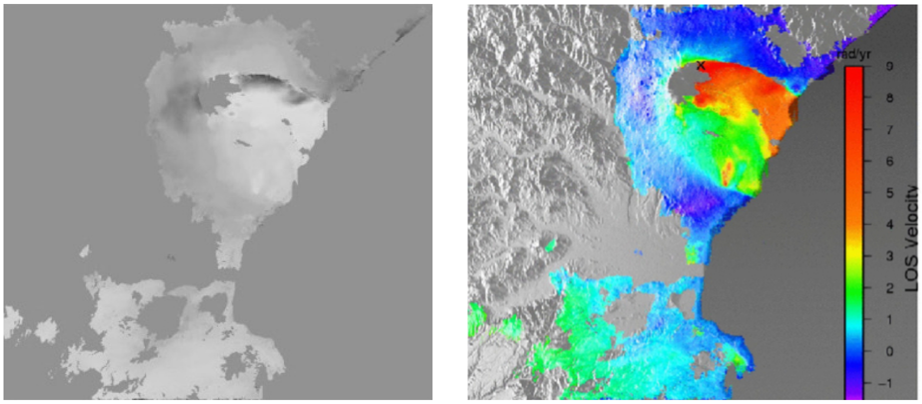

(a) Total phase delays on 2004/08/04 **(b)** Average LOS velocities in rad/year

Fig. 2 The Envisat SITS covering Mount Etna: a) Total phase delays on 2004/08/04, b) Average LOS velocities in rad/year

The quantization of the series was performed as follows. The method does not aim to report fine variations of the pixel values, but is designed to separate, over space and time, relevant phenomena for summarization purpose. Here three intervals (low/medium/high) are sufficient to sketch the global shape of a signal, and thus a simple quantization is made with three symbols (1, 2 and 3) by using the 33rd and the 66th percentiles. The study of the impact of the number of quantization levels made by Julea et al (2011) showed that three symbols are appropriated for GFS-patterns on SITS, however different settings could be needed for other kinds of patterns or other kinds of data. Since pixel values correspond to displacements measured with respect to the same acquisition used as a reference, the quantization is performed over the pixel values of the whole SITS at the same time. The three intervals obtained are: positive values represented by symbol 3 (i.e., motions away from the satellite), negative values close to 0 represented by symbol 2 (i.e., small motions towards the satellite) and the other negative values represented by symbol 1 (i.e., strong motions towards the satellite).

The second SITS is an optical one, built from acquisitions made by Landsat 7. It covers the south-east coast of New Caledonia in the area of Yaté where there is a large nickel mine which is located near a UNESCO World Heritage Site called *the lagoons of New Caledonia*. It contains 16 multispectral images (513x513) acquired between 2000 and 2011, at a spatial resolution of 30 m, and with significant cloud presence on at least half of the series. The images (in a RGB color space) and their dates of acquisition are given in Figure 3. Beside the presence of clouds, this figure shows two additional sources of difficulties for analyzing such a series. Firstly, the dates indicate very different time intervals between the acquisitions, and secondly, some images are impacted by sensor defects and by artifacts (pixels with value of 0, depicted in black). It is important to notice that the summarization method is applied here without any additional processing to handle the clouds, the image quality degradation sources or the irregular time spacing.

This SITS is targeted towards environmental studies, and its core part is a synthetic band giving the *Normalized Difference Vegetation Index (NDVI).* This index (Lillesand et al, 2014) has been designed to express the presence of biomass and is computed using the Red (R) and the Near Infra-Red (NIR) bands as follows: $NDVI = (NIR-R)/(NIR+R)$. The NDVI values vary between -1.0 (absence of biomass) to +1.0 (strong presence of biomass). The summarization method is applied to this synthetic band, and as for the previous dataset only the global shape of the signal is retained, using here a simple quantization of the NDVI values into 3 levels based on the 33^{rd} and the 66^{th} percentiles. Low NDVI index values are denoted by symbol 1, medium ones by symbol 2, and the high ones by symbol 3. For this optical SITS no correction has been made on atmospheric and lightning conditions, thus the intensities in an image cannot be easily compared to intensities in other images. Hence, for such a SITS, the quantization is performed on each image separately.

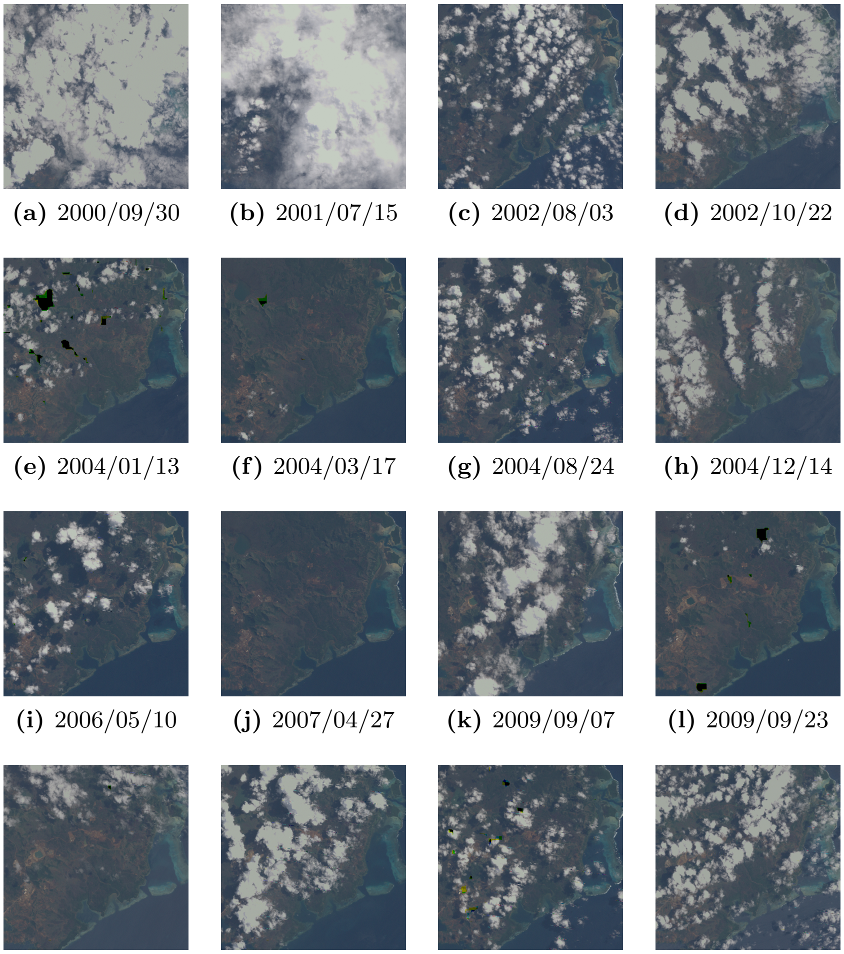

(a) 2000/09/30 (b) 2001/07/15 (c) 2002/08/03 (d) 2002/10/22

(e) 2004/01/13 (f) 2004/03/17 (g) 2004/08/24 (h) 2004/12/14

(i) 2006/05/10 (j) 2007/04/27 (k) 2009/09/07 (l) 2009/09/23

(m) 2009/10/09 (n) 2010/03/02 (o) 2010/04/03 (p) 2011/02/01

Fig. 3 RBG composite images of the Landsat 7 SITS - ©USGS/NASA 2000 - 2011, LP DAAC distribution

4.2 Parameter settings

Since phenomena of interest captured by a SITS are physical phenomena, they are likely to exhibit at least a limited degree of spatial autocorrelation. Thus, to ensure that the regularity depicted by a CE-map is spatially coherent and cannot occur only in rather isolated locations, we required that on average a pattern occurrence have more than half of its 8 nearest neighbors exhibiting the same evolution pattern over time. This is obtained simply by using a minimum average connectivity threshold κ equals to 5. This setting is very general and is used for both series.

To build the SITS summary by taking into account the largest collection of CE-maps, the minimum support threshold σ (i.e. the minimum covered area) is simply set so as to retrieve the greatest number of CE-maps. Figure 4 shows, for each dataset, the number of CE-maps obtained for σ ranging from less than 1% to more than 10% of the surface of an image. According to these measures, σ is set to 7000 for both datasets, corresponding respectively to 508 CE-maps for the Envisat SITS and 297 CE-maps for the Landsat 7 SITS.

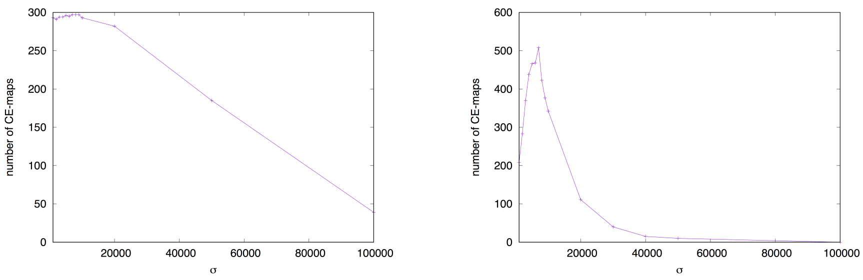

(a) Landsat 7 SITS

(b) Envisat SITS

Fig. 4 Number of CE-maps vs. minimum support threshold σ

Finally, the last parameter to be set is the number of random swap attempts made by the CE-map ranking process. As described in Section 3, based on the current empirical recommendations, the number of random swap attempts used in the summarization method is chosen to be about 20 times the size of the dataset. For the Envisat SITS, we have $553 \times 598 \times 16 = 5\,291\,104$ pixel values, and for the Landsat 7 SITS, we have $513 \times 513 \times 16 = 4\,210\,704$ values. Thus, for both datasets, the number of random swap attempts is set to 100 000 000, which is approximately equal to 20 times the number of pixel values.

Figures 5 and 6 show that such a mixing of the data is sufficient in order to apply a selection by a top-k highest and top-k lowest strategy. In these figures, for each CE-map ranked in the top-20 highest or top-20 lowest NMI values (rank using 100 000 000 swap attempts), a curve is plotted to present the variation of the NMI of this map according to different numbers of swap attempts. Indeed, the ranking is not likely to remain constant because of the intrinsic stochasticity of the process, but the figures show that a mixing of 100 000 000 swap attempts is sufficient to obtain a rather stable ranking. These results were obtained for series encoded with 3 symbols. Figure 7 and Figure 8 show the NMI variations of the top-20 highest and the top-20 lowest CE-maps obtained when encoding the same SITS with 6 symbols. As can be observed, for these datasets, a size of $|\mathcal{D}| = 6$ only has a slight impact on mixing times and setting the number of swap attempts to 20 times the size of the dataset would remain a safe choice. However, as aforementioned, there is no conclusive theoretical result in the literature for setting the appropriate number of swap attempts. Thus, if a larger number of symbols is used, it would be safer to monitor the ranking stability so as to ensure that the number of swaps is sufficient.

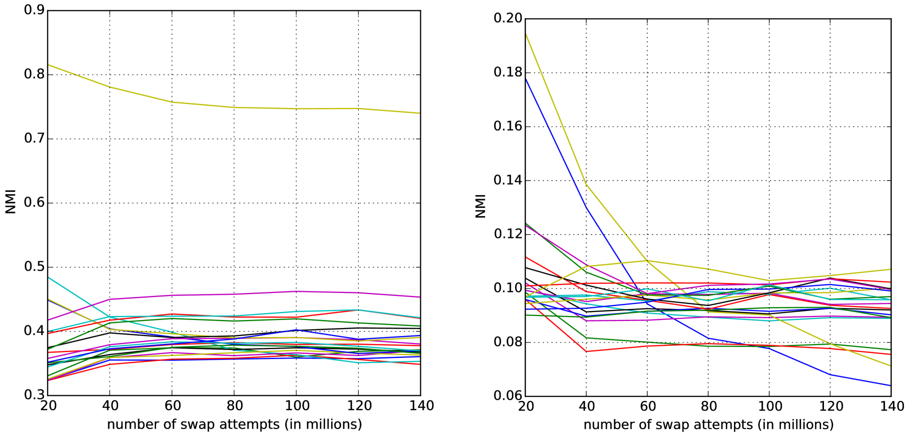

(a) NMI variation of the top-20 highest rank CE-maps

(b) NMI variation of the top-20 lowest rank CE-maps

Fig. 5 NMI vs. number of swap attempts for CE-maps over the Envisat SITS

Another stability aspect is the stability of the ranking with respect to different randomizations with the same number of swap attempts. For each of the two SITS, 1000 swap randomized datasets have been built (for 100 000 000 swap attempts), and the ranks of all CE-maps have been computed for these 1000 randomizations. Figure 9 shows for each CE-maps the mean and the standard deviation of its rank. It can be observed that the proposed mixing strategy yields rankings whose low ranks and high ranks are rather stable with standard deviations of about 1. As the maps belonging to both ends of the ranking are the one selected by the method, this behavior advocates for the use of one single swap randomized dataset, which is of course more efficient than relying on the

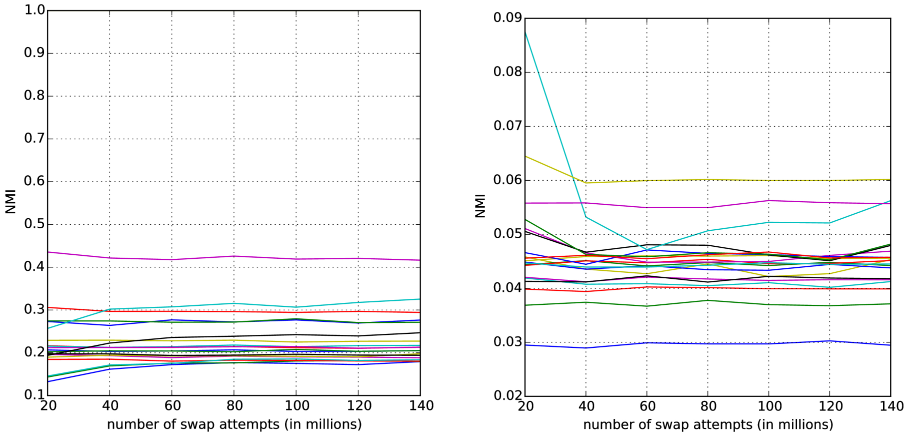

(a) NMI variation of the top-20 highest rank CE-maps

(b) NMI variation of the top-20 lowest rank CE-maps

Fig. 6 NMI vs. number of swap attempts for CE-maps over the Landsat 7 SITS

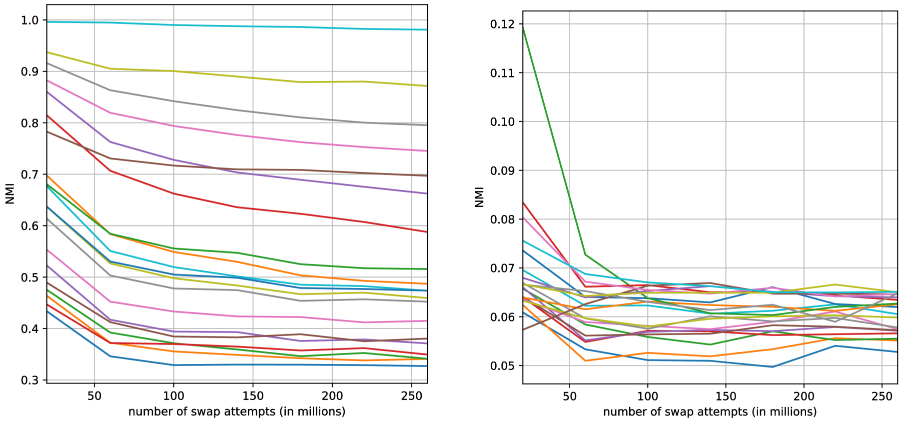

(a) NMI variation of the top-20 highest rank CE-maps, $|\mathcal{D}| = 6$

(b) NMI variation of the top-20 lowest rank CE-maps, $|\mathcal{D}| = 6$

Fig. 7 NMI vs. number of swap attempts for CE-maps over the Envisat SITS encoded with $|\mathcal{D}| = 6$ symbols

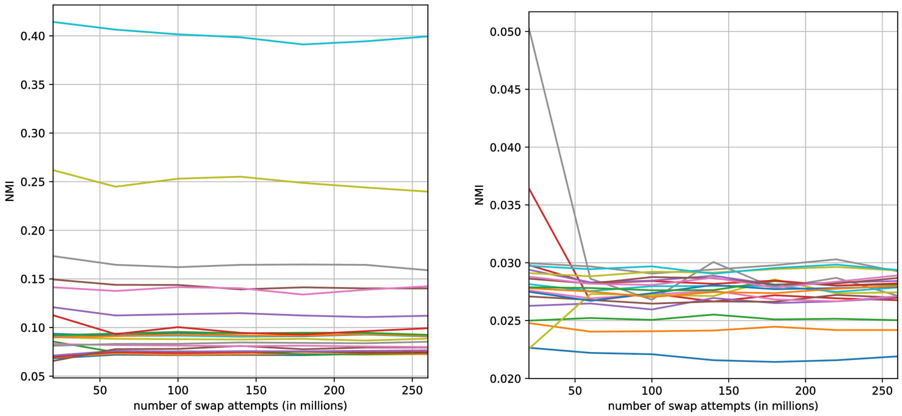

(a) NMI variation of the top-20 highest rank CE-maps, $|\mathcal{D}| = 6$.

(b) NMI variation of the top-20 lowest rank CE-maps, $|\mathcal{D}| = 6$.

Fig. 8 NMI vs. number of swap attempts for CE-maps over the Landsat SITS encoded with $|\mathcal{D}| = 6$ symbols

computation of an empirical p-value for which hundreds of randomized datasets are needed.

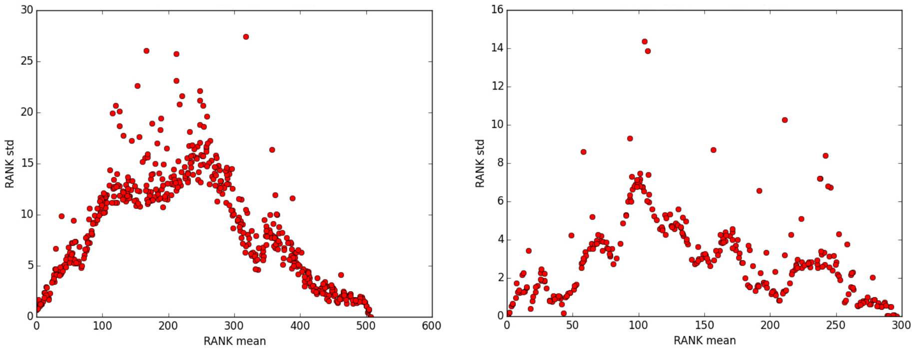

(a) Landsat 7 SITS

(b) Envisat SITS

Fig. 9 Ranking stability: rank std. vs. rank mean

4.3 Results and discussion

The different steps of the summarization method were implemented in C and Python and executed on a standard computing platform (Intel Core i7, 2.8 GHz), using a single-threaded implementation. These software modules are part of the

SITS-P2miner tool (Nguyen et al, 2016) that can be freely downloaded (SITS-Miner-team, 2016).

The whole process includes the image series quantization, the GFS-pattern extraction, the CE-map computation, the randomization and the CE-map ranking. It used no more than 700 MB of RAM, and terminates in less than 380 seconds for the Envisat SITS and in less than 320 seconds for the Landsat 7 SITS.

Three steps within the process are simple to perform: the quantization, the CE-map computation and their ranking. For the first one, the quantization of the series, the cost is proportional to the size of the SITS. If the dataset is already randomized, then the operations corresponding to the drawing and to the NMI computation for all maps, have a cost proportional to the number of maximal GFS-patterns. Since the maps can be processed independently this can be done concurrently if needed. The last operation, i.e., obtaining the final ranking, then consists in sorting the maps according to the NMI values.

Regarding the cost of the two other more elaborate steps of the method, the GFS-pattern extraction and the randomization itself, the execution times have been measured when increasing the sizes of the two SITS from 16 to 26 images. These series have been built by picking at random the images of the original SITS without using the same image more than twice. For a sequence of size n, the number of possible subsequences grows exponentially with n and this effect is coherent with the increase of the GFS-pattern extraction time reported in Figure 10. For the randomization, the number of swap attempts has been increased proportionally to the number of images (starting from 100 000 000) to account for the recommended parameter setting. Figure 11 shows that, as can be expected, the cost of this step increases linearly with respect to the number of images. As will be shown in the next sections, the length of the SITS used in the experiments is sufficient in the context of decadal phenomena of crustal deformation and environmental changes. However, if significantly larger series are required, then the limiting step of the method is clearly the computation of the GFS-patterns. A possible way to tackle such datasets is based on the following remark. The method only needs the maximal patterns, and the maximal sequential patterns satisfying a minimum support constraint are a subset of the *closed* sequential pattern, as defined by Yan et al (2003), satisfying the same constraint. A significant speed-up could then be expected by adapting the so-called *BIDE* algorithm designed by Wang and Han (2004) to extract frequent closed sequential patterns with runtimes that are several orders of magnitude lower than the time needed to compute all the frequent patterns. In addition, even more rapid closed pattern extractions could be achieved by a parallel mining approach such as the one of Yu et al (2012).

In the following sections, the results obtained on the Envisat SITS and on the Landsat 7 SITS are presented and discussed with respect to the phenomena reported in the literature. For each SITS, six CE-maps are given: three different CE-maps having the highest NMI values, and three different ones having the lowest NMI. Of course, more maps could be shown to a user who would like to

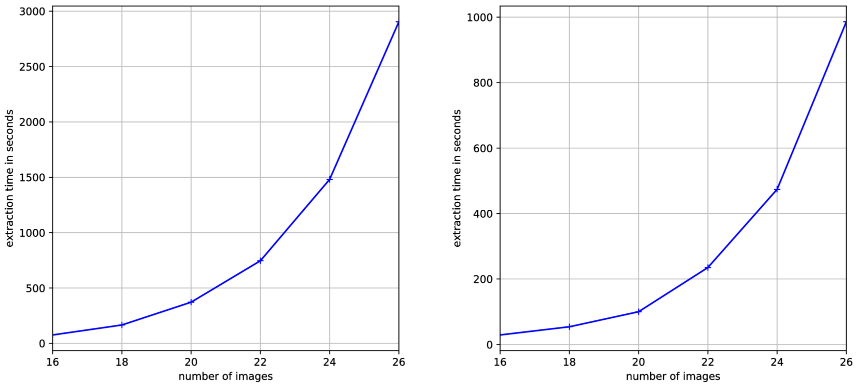

(a) Landsat 7 SITS

(b) Envisat SITS

Fig. 10 Scalability: extraction times vs. SITS sizes

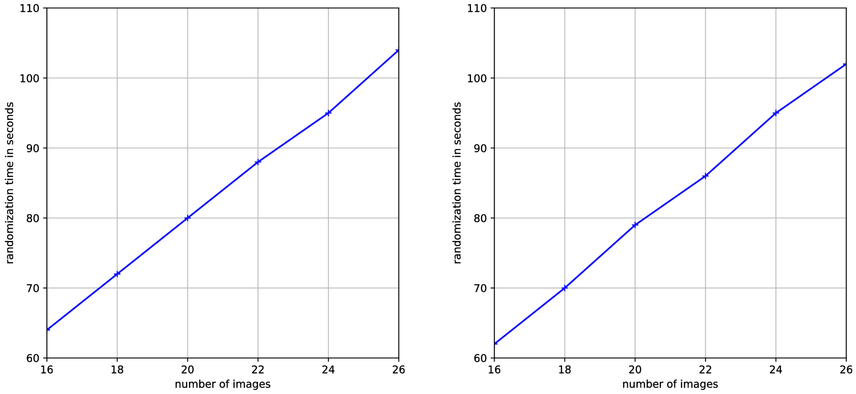

(a) Landsat 7 SITS

(b) Envisat SITS

Fig. 11 Scalability: randomization times vs. SITS sizes see a more detailed summary. The color scale used for all CE-maps is the one given in Figure 1e.

Summarization of the Envisat SITS for crustal deformation monitoring Figure 12 presents top-ranked CE-maps obtained for the first dataset (DEM used as background image). The acquisitions were made during ascending tracks parallel to the left border of the maps, the satellite being located outside of the maps on the left, and looking toward the right. The direction of the north is indicated by an arrow on the first map Figure 12a.

The deformation in this area is known to be characterized by two main domains as explained by Bonforte et al (2011): one on the western and northern flanks of the volcano, which is subject to a global uplift and a roughly radial deformation (westward and northward); and another one combining lateral eastward and vertical downward motions and that is located on the eastern and southeastern flanks. The first one is reflected through the majority of 1 (motion towards the satellite) in the evolution patterns associated with the maps of Figures 12a, 12c and 12f. The second is underlined by the majority of 3 (motion away from the satellite) in the pattern whose map is shown in Figure 12d.

Besides these rather regular deformations, two other more variable phenomena operate on the area, one due to the inflation/deflation phases of the volcano and the other one being recorded as local changes of the relative magnitude of the motions along the three spatial dimensions. The inflation/deflation phases lead to an elastic response observed on the western and southern flanks (Bonaccorso et al, 2006). At some dates, this phenomenon could counterbalance and even have a greater impact than the global regular deformations, and could then be reflected by the few symbols denoting levels 2 and 3 in the evolution patterns of the CE-maps presented in Figures 12a and 12c. The relative variation of the ground motion according to the three axes south/north, east/west, up/down has been exhibited by the analysis of the measures collected by the Etna GPS station network as reported by Azzaro et al (2013). This phenomenon observed in particular on the southern flank could be related to the smooth changes of displacement along the satellite line of sight that are depicted by transitions between symbols 1, 2 and 3 in the evolution patterns of the maps 12b and 12e.

The maps of the summary cover spatially nearly the whole flank of Mount Etna and highlight coherent areas from a geological/geophysical point of view, as detailed more precisely hereafter. The map in Figure 12a is the one having the lowest NMI. It clearly underlines in the S-SW sector of Mount Etna, on the lower part of the flank, the tholeiitic lava plateau extending from Bronte to Paterno that is described by Branca et al (2008) and Bellotti et al (2010).

The second lowest NMI corresponds to the map of Figure 12b. It depicts phenomena in two main areas: one around the volcano and one at the bottom of the map. Around the volcano, it highlights the S-SW upper part of the flank of the volcano corresponding to the area between the Trecastagni fault and the Mascalucia-Tremestrieri fault. From there, the SE thin line going toward the sea lies along the Aci Trezza fault, while the thin line on the northern flank sketches

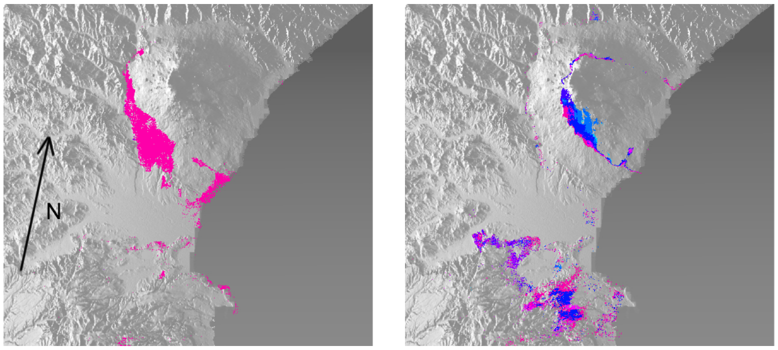

(a) 1^{st} low, $1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 3$

(b) 2^{nd} low, $1 \to 2 \to 2 \to 2 \to 2 \to 2 \to 2 \to 3 \to 3$

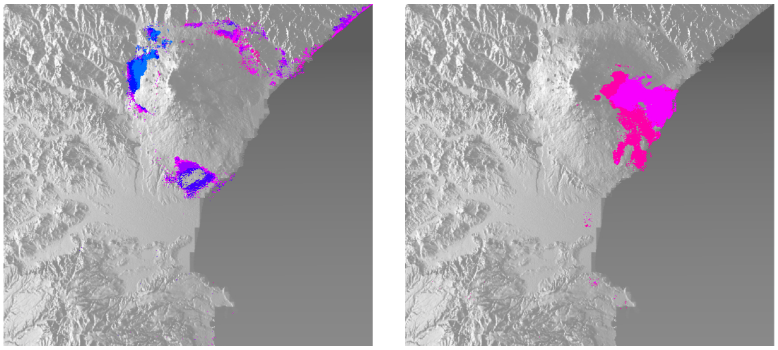

(c) 4^{th} low, $1 \to 1 \to 1 \to 2 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$

(d) 1^{st} high, $1 \to 2 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3 \to 3$

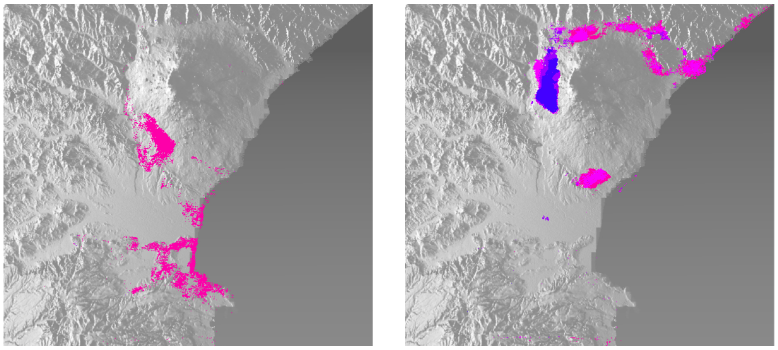

(e) 2^{nd} high, $2 \to 2 \to 1 \to 1 \to 1 \to 2 \to 2 \to 2 \to 2 \to 3$

(f) 8^{th} high, $1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1 \to 1$

Fig. 12 Summary of the InSAR Envisat SITS covering Mount Etna one portion of the Pernicana fault. All these faults are among the main reported ones in the area (Azzaro et al, 2013; Bonforte et al, 2013). The part highlighted at the center of the bottom of the map corresponds to the Villasmundo area where combinations of horizontal movements and subsidence related to groundwater over-exploitation have been observed and reported by Canova et al (2012).

The CE-map having the third lowest NMI is similar to the second one and is not retained. The fourth is different, and is given in Figure 12c. On the western flank and on the northern flank, the map outlines respectively the Ragalna fault system and the Pernicana fault system described by Neri et al (2009). It also sketches the Taormina fault zone (also called Messina-Giardini fault zone) reported by Catalano and De Guidi (2003) in the northeastern area along the coast. The last area exhibited by the map is located S-SE of the volcano, and corresponds to the sides of the anticline termed Misterbianco ridge (between Catania and Misterbianco towns). The strong presence of symbol 1 in the evolution associated to the map, is coherent with the uplift reported for this active anticline by Bonforte et al (2011).

The summary is then completed with three maps taken among the ones having the highest NMI. The first one is given in Figure 12d and, as already mentioned, depicts one of the main domains of deformation surrounding Mount Etna that corresponds to a global eastward and downward motion on the S-SE flank. It can be noticed that there are two parts highlighted by a slight change in the color, from magenta (eastern part) to pink (southeastern part). This indicates that the evolutions are not synchronized on the whole map. Indeed, these two parts are coherent with the two blocks having different velocity profiles identified by Bonforte et al (2011).

The map having the second highest NMI is shown Figure 12e. On the S-SW flank of Mount Etna it completes the map of Figure 12b by underlying a lower part of the flank that corresponds to a block identified as having a different behavior than the upper part (Azzaro et al, 2013). At the bottom, the map essentially sketches faults, in particular the Pedaggagi-Agnone system at the lowest part of the zone and that is presented in the paper of Catalano et al (2010).

From the third to the seventh highest NMI, the maps are very similar to the second one. Indeed, this kind of redundant CE-map corresponds to very similar evolution patterns, and can be filtered out. The next different map is the eighth one and is presented in Figure 12f. It essentially completes the regions highlighted by the map of Figure 12c in particular along the coast for the Taormina fault zone, but also in the north area of the Pernicana fault (northern of Etna), and very clearly for the Misterbianco ridge (southern of Etna) where it exhibits the central upper part of the ridge.

Summarization of the Landsat 7 SITS for environmental monitoring For this SITS, the presented summary consists also of six CE-maps (three different among the highest and among the lowest NMI values). In addition to the references cited hereafter, the interpretation of the summary content was made during two field trips and making use of the knowledge of the area provided by the *Qëhnelö^{TM} environmental management platform* (see Bluecham SAS, 2016).

The summary is given in Figure 13, where the CE-maps are drawn on a RGB composite background image. Here, only one CE-map was repeated. It is the one having the second lowest NMI shown Figure 13b. It reappears for similar patterns as the maps having the third, fourth and fifth lowest NMI, which were then not retained. Indeed, this map of Figure 13b reflects one of the most salient phenomenon of the area, that is the presence of the so called *maquis* on ultramafic rocks in New Caledonia (Proctor, 2003). Its biomass variation over time is reflected by the NDVI variation between medium and high levels in the associated evolution pattern.

This map is spatially completed by the one of Figure 13d, which has the highest NMI and that underlines the area corresponding to the ocean itself, for which the presence of water leads to low NDVI values (Lillesand et al, 2014), as captured by the associated evolution pattern. Notice that the map also covers the fringing reefs and their areas of shallow water along the coast that can be seen on the RGB background of the other CE-maps of Figure 13. These reefs are reported on the *Grand Sud* map of Sevin et al (2012). On this CE-map Figure 13d, especially in the bottom right corner, oblique streaks are visible, formed by thin lines of dots where there is no occurrence of the evolution pattern. These streaks correspond to the Landsat 7 defects occurring in the acquisitions from 2003 to the end of the SITS, and due to the failure of the *scan line corrector* of the satellite on May 31, 2003 reported by USGS et al (2003). After 2003, the impact of this failure has been minimized by additional image processing. In the case of the acquisitions forming the SITS used here, the additional processing made was the recovery method presented by Chen et al (2011). However, some defects are still present together with artifacts, as shown for instance in Figure 14a that reproduces the bottom right corner of the area as it appears in the 2011/02/01 acquisition. Notice also that since the acquisitions are not made each time from the same location, then the location of the streaks changes over the series, leading to the irregular footprint of the defects capture by the map Figure 13d. There are other larger parts over the ocean that are not covered by the associated evolution pattern. This is due to the presence in the SITS of two images nearly completely filled with clouds. Indeed, clouds are likely to have low NDVI values, but since here for an optical SITS the quantization retained is performed for each image separately, a cloudy area may have a high NDVI value when compared to the rest of the image, and thus will not be encoded as symbol 1. For instance, on the NDVI image 2001/07/15 shown Figure 14b, this is the case of the bottom right cloud shape (circled). Here, missing occurrences of symbol 1 lead to a footprint having the same shape, and exhibited as an non-covered area in the CE-map of Figure 13d.

Between the areas covered by the maps of Figures 13b and 13d lies the area highlighted by the map given in Figure 13e (second highest NMI) for which the associated evolution indicates a high NDVI value. Indeed, this area corresponds to a vegetation different from the one of Figure 13b, and includes mainly ever-

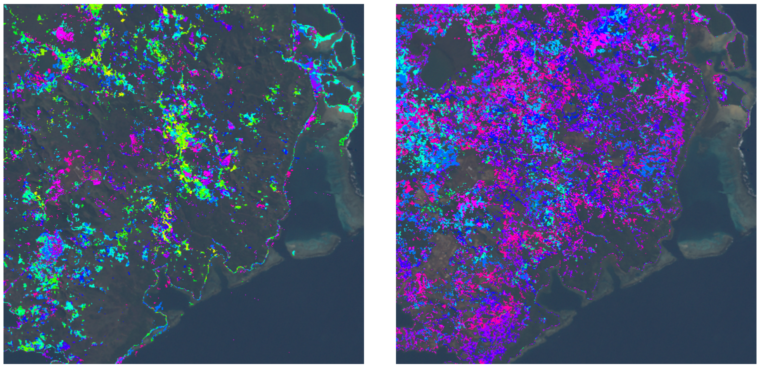

(a) 1st low, $1 \rightarrow 1 \rightarrow 2 \rightarrow 2$

(b) 2nd low, $2 \rightarrow 2 \rightarrow 3 \rightarrow 2 \rightarrow 2 \rightarrow 2 \rightarrow 3$

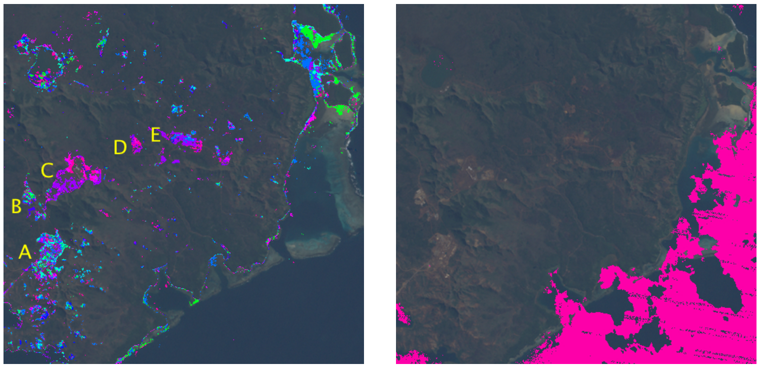

(c) 6th low, $2 \rightarrow 2 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 2$

(d) 1st high, $1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1$

(e) 2nd high, $3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3 \rightarrow 3$

(f) 3rd high, $1 \rightarrow 2 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 1$

Fig. 13 Summary of the Landsat 7 SITS covering the area of Yaté in New Caledonia (RGB composite background)

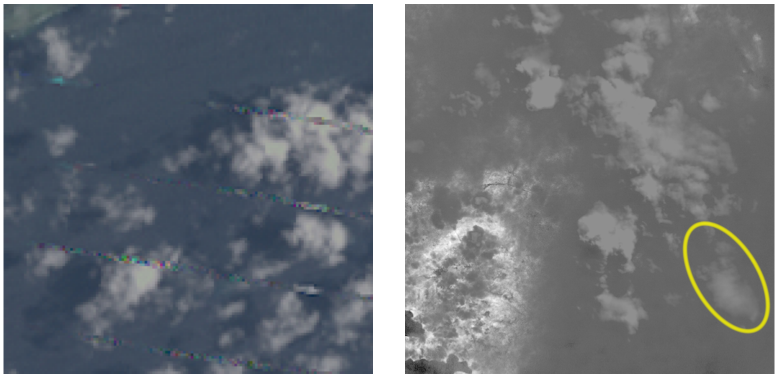

(a) Defect appearing in the bottom right corner of the 2011/02/01 acquisition (RGB composite image)

(b) The 2001/07/15 NDVI image of the SITS. Strongly impacted by the presence of clouds

Fig. 14 Landsat 7 defect (RGB) and cloud aspect (NDVI)

green conifers. On this part of the coast, it consists essentially of endemic Araucaria species, such as the famous Araucaria columnaris as reported by Wilcox and Platt (2002). This kind of vegetation also covers a large part of the *Cap Coronation*, and of the two coastal islands of Nuu (Nou) and Néaé. These three zones are respectively the three shapes underlined in the upper right corner of the map along the coast.

As in Figure 13d, the CE-map having the third highest NMI underlines phenomena related to water and clouds. This map is given in Figure 13f, where the associated pattern indicates that close to the beginning of the SITS, at least in one acquisition the NDVI level was not low (not encoded as symbol 1) but has a medium value (symbol 2). In the upper left part this corresponds to a temporary drying up of a lake. The latter is the main lake of the area covered by the SITS and is visible in Sevin et al (2012). The coastal areas in the upper right and lower left corners point out parts covered by water but that emerge during a certain period (e.g., reef flats). Such water level reductions are likely to be due in these zones to modification of the coastal currents and to changes in the deposition of sediments. The other elements appearing on the map are located in the bottom right. They correspond to borders of cloud slices where the NDVI values were quantified as symbol 2 in an acquisition close to the beginning of the SITS. Indeed, their shapes are consistent with the footprints of the clouds on Figure 13d.

Beyond biomass and water, the NDVI index value can be related to other elements, since it depends on emission/reflection in the red and in the near infra-red bands. In the south-east region of New Caledonia, variations in the red band can be due to the uncovering/covering of the red laterite soil layer, while thermical effects of industrial facilities can cause changes in the near infra-red band. The summary points out such phenomena in Figure 13c. In particular it exhibits the five sites of mining activities reported in the mining cadastre view of Sevin et al (2012) and labeled (on their left) on this CE-map as follows: (A) processing plant, (B) living quarters, (C) storage basin (for solid residual products), (D) preprocessing unit, and (E) mining area. For the storage bassin (C), the color scale denotes that the evolution pattern was found earlier in one part (in purple) than in the other part (in magenta). This is coherent with the development of the mining activity that leads to an extension over time of the storage area towards the north-east.

The CE-map of Figure 13c also highlights two other phenomena in the top right corner. Firstly, it depicts in blue the mouth of the Kuébéni (Kubini) river. This river is known for its load of red clay and corresponding sediment deposition here at its mouth as explained by Bird et al (1984), where the evolution pattern seems to reflect variations in this load and/or deposition. Secondly, it underlines in green several flat coastal areas in shallow water that are subject to silting and water level variations. Another change of water level is reported in the map for a lakeshore in the top left corner (the lake highlighted in the map of Figure 13f). The last important element of Figure 13c is a thin line along the coast of the main land (not encompassing the fringing reefs) and along the coast of the small islands (top right corner). No existing study of coastal changes were found in the literature for this particular area, but coastline modification in New Caledonia is a recent and active subject of investigation with works such those presented by Garcin and Vendé-Leclerc (2014). This coastal line is even more clearly exhibited on the last CE-map to be commented on in this summary and is presented in Figure 13a. Indeed, this map is the best ranked map with respect to the low NMI criterion, and thus is the least likely to occur in swap randomized data. It reports nearly all the phenomena that have been more finely isolated in the map of Figure 13c. The other elements, specific to this map in Figure 13a, are essentially the yellow and green areas in the upper left, bottom left and central parts. This yellow and green zones correspond to areas where vegetation regrowth took place after mine earthworks (e.g., leveling in the center) and after one of the recurrent bushfires (McCoy et al, 1999) such as the fire in the bottom left part on December 13, 2008. It can be noted that the summarization method isolated phenomena not only over space, but also separated them over time, as shown by the central green/yellow locations in the map of Figure 13a that are also covered in the CE-map of Figure 13b by evolutions ending later in the series (colors blue/purple/magenta).

Significance of the results Even though some global trends could be visually detected on SITS, when it comes to consider a set of images and all of its possible subsets, for all possible zones and all possible evolutions, then the task is too complex to be human-tractable. A common approach used by experts is to build a single image where each pixel represents a measure obtained by aggregating the values at the corresponding location over time. This has been done, for example, by Doin et al (2011) to explore the Envisat SITS covering

Mount Etna, resulting in average velocity maps similar to the one shown in Figure 2b in Section 4.1. As can be observed, such an average map exhibits different zones but does not isolate different evolutions over space and time. The CE-maps are thus important refinements that cannot be obtained by inspecting SITS or average maps manually. However, there is one particular case where a CE-map and an average map can be very similar. This arises when a pattern has a size equal to the number of images of the series. In this case, the pixels covered by the pattern all have the same value in the average map. Applying a filter to select these pixels thus gives the same locations as the ones highlighted in the CE-map. In figures 12 and 13, this occurs only for the map of Figure 13d.

As detailed in the previous sections, the CE-maps reported evolutions that are consistent with the current knowledge of the phenomena of the studied areas. Besides refining aggregation-based maps, such as average maps, by a disentanglement of evolutions, they also complete ground-based measurements and field trip observations that only account for a limited number of locations. The CE-maps give insights about the whole spatial extent of evolutions to understand the phenomena, and also, for instance, to decide of future locations of ground measurements/investigations.

With regard to the significance of the selected CE-maps with respect to the SITS itself, the method relies on a randomization approach. Such frameworks have been used to assess the significance of data mining results (e.g., clustering, frequent itemsets) obtained from Boolean matrices (e.g., Gionis et al, 2007) by comparing scores on a dataset with scores on randomized versions of this dataset. In the case of CE-maps, this technique is extended to handle series of images in symbolic SITS. More generally, the comparison of pattern scores between datasets (or parts of a dataset) is a fruitful mining approach that led to the development of three main families of algorithms: contrast set mining (e.g., Bay and Pazzani, 2001), emerging pattern mining (e.g., Dong and Li, 1999) and subgroup discovery (e.g., Klösgen, 1996). As shown by Novak et al (2009), they are all based on pattern occurrence counting, and aim to find descriptions that optimize a trade-off between the amount of data they covered and their ability to precisely characterize specific parts of the data. They are all designed to find descriptions of subsets of the data that are identified in the input by a reference labelling (or partitioning). For CE-maps, there is no targeted partition of the data, but the selection focuses on two kinds of maps: the ones that are easily lost after randomization because they reflect evolutions that are very specific to the original dataset, and the maps that remain largely unchanged after randomization since they report evolutions strongly represented in the SITS. In addition, the ranking does not rely only on measures based on number of occurrences of patterns, but it uses a NMI measure that incorporates the location of the occurrences in space and time.

The method leads to the selection of CE-maps depicting evolutions that are coherent in space and significant with respect to a swap randomization procedure, but the interpretation of an expert is needed to relate these evolutions to real phenomena. And even if the method can be applied on low quality data (e.g., clouds, acquisition defects), it only reports in maps information that is present in the SITS. It is thus recommended that domain experts have some knowledge of the level of preprocessing that has been applied to the SITS so as to avoid any misleading interpretations.

5 Conclusion

The development of acquisition means and of open data policies favors the availability of increasing collections of SITS prepared for various studies, and the reuse of these SITS in other applications is an important opportunity, especially for geosciences. Beyond the use of metadata (e.g., location, geometry, date, frequency band) to search a collection of SITS, the quick browsing through spatiotemporal phenomena captured in the data could provide evidence of the potential interest of some of these SITS for a new study. With this aim, an unsupervised SITS summarization method was proposed in this paper. It consists in building a summary of a SITS as a set of maps depicting *core evolutions* over space and time, where these maps are ranked by comparing them to the maps of the same evolutions in randomized data, using their information content.

In contrast to many existing SITS processing techniques, the approach is not targeted to specific objects/regions or to certain changes, and is general enough to be applied to both radar and optical SITS. This degree of methodological objectivity/neutrality of the method has been observed in the summaries obtained for two real SITS. They indeed reported very different elements and phenomena, e.g., faults, ridges, subsidences, uplifts, anthropic activities, surface earthworks/erosion, shallow water level variations, different kinds of vegetation and vegetation changes. The restriction made, with respect to the phenomena, is that they must be noticeable in a single band and still be present after a simple quantization of the data. However, the method is clearly not intended to compete with techniques targeted at dedicated fine-grained analyses. It aims to ease the selection, among SITS, of the ones from which it could be worth undertaking such more specific and resource consuming processes. The reported experiment showed, for instance, that the approach can exhibit evidence of subsidence as well as modifications in sediment deposition in river mouths or variations of water level on lake shores. Such phenomena may be related or due to anthropic activities (like oil and gas exploitation/exploration) and their observation in a SITS can be of primary interest to trigger deeper studies and drive new adaptations of these anthropic processes.

The produced summaries are easy to read, in the sense that they separate the elements/phenomena over space and/or time, as shown by the interpretation of the map contents made in Section 4. Of course, even if this disentanglement favors the readability, intended users such as geoscientists still need to have some expertise in the family of spatiotemporal phenomena that are exhibited. Besides being applicable to different data (radar and optical), the method has a simple but effective parameter setting procedure and was able to handle common SITS quality problems (clouds, different lightning conditions, acquisition defects, artifacts, irregular acquisition intervals, missing parts in images) without additional processing. This allows the method to offer some exhaustiveness, in the sense that the maps can account for the whole SITS over space and time, without removing images or parts of images because of some of the aforementioned defects. Accordingly, their is no temporal constraint on the phenomena. For example, maps are not limited to phenomena detected between consecutive images or starting on the first image, but can occur anywhere and can spread over short or large time intervals. However the method could be tailored to more specific applications. This can be done, for instance, by relying on preprocessings such as cloud masking or applying syntactic constraints to patterns to focus on specific phenomena and also filtering out patterns that could be redundant, e.g., patterns covering the same geographical zones and/or the same temporal periods.

There are several promising directions for future work. Firstly, with respect to the data encoding itself, could more information be used when ranking the maps? For instance, is it valuable to take into account the ordering of the symbols used for the quantization (e.g., giving less weight to a difference between symbol 1 and symbol 2 and more weight to a difference between symbol 1 and symbol 3)? The second direction would be to reflect in the maps some phenomena that have a smaller extent over space. The method was able to exhibit phenomena that covered at least 2 or 3% of the area, but how could one enable it to focus on phenomena with a much smaller spatial footprint? Could it be possible, in such an unsupervised process, to point out phenomena affecting only a few pixels if there is also a large collection of phenomena spreading over wide areas? A third aspect to investigate would be to quantify the information loss when trying to build a summary. A starting point could be the approaches that select patterns leading to the best compression of a dataset (using a minimum description length criterion) as done by Vreeken et al (2011). However, even if this family of research has been extended to handle sequences of symbols by Tatti and Vreeken (2012) or Lam et al (2014), using it to handle together compression over space and time for spatiotemporal data such as SITS is still an open question.

Acknowledgements

The Lansat 7 SITS was retrieved from the online Data Pool, courtesy of the NASA Land Processes Distributed Active Archive Center (LP DAAC), USGS / Earth Resources Observation and Science (EROS) Center, Sioux Falls, South Dakota `https://lpdaac.usgs.gov/data_access/data_pool`. The authors wish to thank the European Space Agency (ESA) for providing the ENVISAT SAR data over Mount Etna, and the Yaté rural district of New Caledonia for its support.

Bibliography

Aggarwal CC, Hinneburg A, Keim DA (2001) On the surprising behavior of distance metrics in high dimensional space. In: Proceedings of the 8th International Conference on Database Theory (ICDT'01), London, UK, pp 420–434

Agrawal R, Srikant R (1995) Mining sequential patterns. In: Proceedings of the 11th IEEE International Conference on Data Engineering (ICDE'95), Taipei, Taiwan, pp 3–14

Akbari V, Doulgeris AP, Eltoft T (2014) Monitoring glacier changes using multitemporal multipolarization SAR images. IEEE Transactions on Geoscience and Remote Sensing 52(6):3729–3741

Alatrista Salas H, Bringay S, Flouvat F, Selmaoui-Folcher N, Teisseire M (2012) The pattern next door: Towards spatio-sequential pattern discovery. In: Proceeding of the 16th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'12), Kuala Lumpur, Malaysia, pp 157–168

Alonso-Gonzalez A, Lopez-Martinez C, Salembier P (2012) Filtering and segmentation of polarimetric sar data based on binary partition trees. IEEE Transactions on Geoscience and Remote Sensing 50(2):593–605

Amitrano D, Ciervo F, Di Bianco P, Di Martino G, Iodice A, Mitidieri F, Riccio D, Ruello G, Papa MN, Koussoube Y (2015) Monitoring soil erosion and reservoir sedimentation in semi-arid region through remote sensed SAR data: a case study in Yatenga Region, Burkina Faso. In: Proceedings of the 12th International Association for Engineering Geology and the Environment Congress, New Delhi, India, vol 3, pp 539–542

Azzaro R, Bonforte A, Branca S, Guglielmino F (2013) Geometry and kinematics of the fault systems controlling the unstable flank of Etna volcano (Sicily). Journal of Volcanology and Geothermal Research 251:5–15

Bay SD, Pazzani MJ (2001) Detecting group differences: Mining contrast sets. Data Mining and Knowledge Discovery 5(3):213–246

Bellotti F, Branca S, Groppelli G (2010) Geological map of Mount Etna West Rift (Italy). Journal of Maps 6(1):96–122

Bird EC, Dubois JP, Iltis JA (1984) The impacts of opencast mining on the rivers and coasts of New Caledonia. Tech. Rep. NRTS-25/UNUP-505, p. 64, United Nations University

Bluecham SAS (2016) Qëhnelö$^{\mathrm{TM}}$ Environmental Management Platform, full Yaté area version at `http://www.yate.nc/` (accessed 08/25/2018)

Bonaccorso A, Bonforte A, Guglielmino F, Palano M, Puglisi G (2006) Composite ground deformation pattern forerunning the 2004-2005 Mount Etna eruption. Journal of Geophysical Research: Solid Earth 111(B12):1–11

Bonforte A, Guglielmino F, Coltelli M, Ferretti A, Puglisi G (2011) Structural assessment of Mount Etna volcano from permanent scatterers analysis. Geochemistry, Geophysics, Geosystems 12(2):1–19

Bonforte A, Federico C, Giammanco S, Guglielmino F, Liuzzo M, Neri M (2013) Soil gases and SAR measurements reveal hidden faults on the sliding flank of

Mount Etna (Italy). Journal of Volcanology and Geothermal Research 251:27 – 40

Bontemps S, Bogaert P, Titeux N, Defourny P (2008) An object-based change detection method accounting for temporal dependences in time series with medium to coarse spatial resolution. Remote Sensing of Environment 112:3181–3191

Branca S, Coltelli M, De Beni E, Wijbrans J (2008) Geological evolution of Mount Etna volcano (Italy) from earliest products until the first central volcanism (between 500 and 100 ka ago) inferred from geochronological and stratigraphic data. International Journal of Earth Sciences 97(1):135–152

Burdick D, Calimlim M, Flannick J, Gehrke J, Yiu T (2005) Mafia: a maximal frequent itemset algorithm. IEEE Transactions on Knowledge and Data Engineering 17(11):1490–1504

Canova F, Tolomei C, Salvi S, Toscani G, Seno S (2012) Land subsidence along the Ionian coast of SE Sicily (Italy), detection and analysis via Small Baseline Subset (SBAS) multitemporal differential SAR interferometry. Earth Surface Processes and Landforms 37(3):273–286

Cao H, Mamoulis N, Cheung DW (2005) Mining frequent spatio-temporal sequential patterns. In: Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM'05), Washington, DC, USA, pp 82–89

Cao H, Mamoulis N, Cheung DW (2007) Discovery of periodic patterns in spatiotemporal sequences. IEEE Transaction on Knowledge and Data Engineering 19(4):453–467

Carvalho DFd, Durigon VL, Antunes MAH, Almeida WSd, Oliveira PTSd (2014) Predicting soil erosion using Rusle and NDVI time series from TM Landsat 5. Pesquisa Agropecuária Brasileira 49:215 – 224

Catalano S, De Guidi G (2003) Late Quaternary uplift of northeastern Sicily: relation with the active normal faulting deformation. Journal of Geodynamics 36(4):445 – 467

Catalano S, Romagnoli G, Tortorici G (2010) Kinematics and dynamics of the Late Quaternary rift-flank deformation in the Hyblean Plateau (SE Sicily). Tectonophysics 486(14):1 – 14

Cauwels P, Pestalozzi N, Sornette D (2014) Dynamics and spatial distribution of global nighttime lights. EPJ Data Science 3(1):1–26

Chen HM, Varshney P, Arora M (2003) Performance of mutual information similarity measure for registration of multitemporal remote sensing images. IEEE Transactions on Geoscience and Remote Sensing 41(11):2445–2454

Chen J, Zhu X, Vogelmann JE, Gao F, Jin S (2011) A simple and effective method for filling gaps in Landsat ETM+ SLC-off images. Remote Sensing of Environment 115(4):1053 – 1064

Cobb GW, Chen YP (2003) An application of markov chain monte carlo to community ecology. The American Mathematical Monthly 110(4):pp. 265–288

Collignon A, Maes F, Delaere D, Vandermeulen D, Suetens P, Marchal G (1995) Automated multi-modality image registration based on information theory. In: Proceedings of the 14th International Conference on Information Processing in Medical Imaging (ICIP'95), Brest, France, pp 263–274

Coppin P, Jonckheere I, Nackaerts K, Muys B, Lambin E (2004) Digital change detection methods in ecosystem monitoring: a review. International Journal of Remote Sensing 25(9):1565–1596

Cover T, Thomas J (1991) Elements of information theory. Wiley-Interscience, New York, NY, USA

Crawford CJ, Manson SM, Bauer ME, Hall DK (2013) Multitemporal snow cover mapping in mountainous terrain for landsat climate data record development. Remote Sensing of Environment 135:224 – 233

Dogan O, Perissin D (2014) Detection of multitransition abrupt changes in multitemporal sar images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7(8):3239–3247

Doin MP, Lasserre C, Peltzer G, Cavalié O, Doubre C (2009) Corrections of stratified tropospheric delays in SAR interferometry: Validation with global atmospheric models. Journal of Applied Geophysics 69(1):35–50

Doin MP, Lodge F, Guillaso S, Jolivet R, Lasserre C, Ducret G, Grandin R, Pathier E, Pinel V (2011) Presentation of the small baseline NSBAS processing chain on a case example: the Etna deformation monitoring from 2003 to 2010 using Envisat data. In: Proceedings of the European Space Agency Workshop on Advances in the Science and Applications of SAR interferometry Fringe (Fringe'11), Frascati, Italy, pp 3434–3437

Doin MP, Twardzik C, Ducret G, Lasserre C, Guillaso S, Jianbao S (2015) InSAR measurement of the deformation around Siling Co Lake: Inferences on the lower crust viscosity in central Tibet. Journal of Geophysical Research: Solid Earth 120(7):5290–5310

Dong G, Li J (1999) Efficient mining of emerging patterns: Discovering trends and differences. In: Proceedings of the 5th International Conference on Knowledge Discovery and Data Mining (KDD'99), San Diego, California, USA, pp 43–52

Duede E, Zhorin V (2016) Convergence of economic growth and the great recession as seen from a celestial observatory. EPJ Data Science 5(1):1–29

Fahnestock M, Scambos T, Moon T, Gardner A, Haran T, Klinger M (2016) Rapid large-area mapping of ice flow using landsat 8. Remote Sensing of Environment 185:84 – 94

Gallucio L, Michel O, Comon P (2008) Unsupervised clustering on multi-components datasets: Applications on images and astrophysics data. In: 16th European Signal Processing Conference (EUSIPCO'08), Lausanne, Switzerland, pp 25–29

Garcin M, Vendé-Leclerc M (2014) Observatoire du littoral de Nouvelle-Calédonie : observations, état des lieux et constats. Tech. Rep. BRGM/RP-63235-FR, p. 125, Bureau de Recherches Géologiques et Minières (BRGM), Nouméa

Gionis A, Mannila H, Mielikäinen T, Tsaparas P (2007) Assessing data mining results via swap randomization. ACM Transactions on Knowledge Discovery from Data 1(3):1–32

Gonçalves G, Duro N, Sousa E, Pinto L, Figueiredo I (2014) Detecting changes on coastal primary sand dunes using multi-temporal Landsat imagery. In: Proceedings of 20th SPIE International Conference of Image and Signal Processing for Remote Sensing, Amsterdam, Netherlands, vol 9244, p 8

Gouda K, Zaki MJ (2001) Efficiently mining maximal frequent itemsets. In: Proceedings of the 2001 IEEE International Conference on Data Mining (ICDM '01), Washington, DC, USA, pp 163–170

Griffith DA, Chun Y (2016) Spatial autocorrelation and uncertainty associated with remotely-sensed data. Remote Sensing, article number 535, 8(7)

Gudmundsson J, Kreveld M, Speckmann B (2007) Efficient detection of patterns in 2D trajectories of moving points. Geoinformatica 11(2):195–215

Gueguen L, Datcu M (2007) Image time-series data mining based on the information-bottleneck principle. IEEE Transactions on Geoscience and Remote Sensing 45(4):827–838

Héas P, Datcu M (2005) Modeling trajectory of dynamic clusters in image time-series for spatio-temporal reasoning. IEEE Transactions on Geoscience and Remote Sensing 43(7):1635– 1647

Huang Y, Zhang L, Zhang P (2008) A framework for mining sequential patterns from spatio-temporal event data sets. IEEE Transactions on Knowledge and Data Engineering 20(4):433–448

Ilsever M, Ünsalan C (2012) Texture analysis based change detection methods. In: Two-Dimensional Change Detection Methods: Remote Sensing Applications, Springer, London, chap 4, pp 35–39

Inglada J, Giros A (2004) On the possibility of automatic multisensor image registration. IEEE Transactions on Geoscience and Remote Sensing 42(10):2104–2120

Inglada J, Favard JC, Yesou H, Clandillon S, Bestault C (2003) Lava flow mapping during the Nyiragongo January, 2002 eruption over the city of Goma (D.R. Congo) in the frame of the international charter space and major disasters. In: Proceedings of the IEEE International Conference on Geoscience And Remote Sensing (IGARSS'03), Toulouse, France, vol 3, pp 1540–1542

Jaccard P (1902) Lois de distribution florale dans la zone alpine. Bulletin de la Société Vaudoise des Sciences Naturelles 38:69–130

Julea A, Méger N, Bolon P, Rigotti C, Doin MP, Lasserre C, Trouvé E, Lăzărescu V (2011) Unsupervised spatiotemporal mining of satellite image time series using grouped frequent sequential patterns. IEEE Transactions on Geoscience and Remote Sensing 49(4):1417–1430

Kayastha N, Thomas V, Galbraith J, Banskota A (2012) Monitoring wetland change using inter-annual Landsat time-series data. Wetlands 32(6):1149–1162

Klösgen W (1996) Explora: A multipattern and multistrategy discovery assistant. In: Advances in Knowledge Discovery and Data Mining, AAAI, pp 249–271

Konings A, McColl K, Piles M, Entekhabi D (2015) How many parameters can be maximally estimated from a set of measurements? IEEE Geoscience and Remote Sensing Letters 12(5):1081–1085

Krylov VA, Moser G, Serpico SB, Zerubia J (2013) False discovery rate approach to image change detection. In: Proceedings of the 2013 IEEE International Conference on Image Processing (ICIP'13), Melbourne, Australia, pp 3820–3824

Lam HT, Mörchen F, Fradkin D, Calders T (2014) Mining compressing sequential patterns. Statistical Analysis and Data Mining 7(1):34–52

Li L, Leung M (2001) Robust change detection by fusing intensity and texture differences. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'01), Kauai, HI, USA, vol 1, pp 777–784

Lillesand T, Kiefer RW, Chipman J (2014) Remote sensing and image interpretation, 7th edn. John Wiley & Sons, New York

Liu Z, He C, Zhang Q, Huang Q, Yang Y (2012) Extracting the dynamics of urban expansion in China using DMSP-OLS nighttime light data from 1992 to 2008. Landscape and Urban Planning 106(1):62 – 72

Lu D, Mausel P, Brondizio E, Moran E (2004) Change detection techniques. International Journal of Remote Sensing 25(12):2365–2407

Lu M, Chen J, Tang H, Rao Y, Yang P, Wu W (2016) Land cover change detection by integrating object-based data blending model of Landsat and MODIS. Remote Sensing of Environment 184:374 – 386

Luo C, Chung SM (2005) Efficient mining of maximal sequential patterns using multiple samples. In: Proceedings of the 2005 SIAM International Conference on Data Mining (ICDM'05), Newport Beach, CA, USA, pp 415–426

Mannila H, Toivonen H, Verkamo AI (1997) Discovery of frequent episodes in event sequences. Data Mining and Knowledge Discovery 1(3):259–289

Marin C, Bovolo F, Bruzzone L (2015a) Building change detection in multitemporal very high resolution sar images. IEEE Transactions on Geoscience and Remote Sensing 53(5):2664–2682

Marin C, Bovolo F, Bruzzone L (2015b) Building change detection in multitemporal very high resolution sar images. IEEE Transactions on Geoscience and Remote Sensing 53(5):2664–2682

McCoy S, Jaffre T, Rigault F, Ash JE (1999) Fire and succession in the ultramafic maquis of New Caledonia. Journal of Biogeography 26(3):579–594

Méger N, Rigotti C, Pothier C (2015) Swap randomization of bases of sequences for mining satellite image times series. In: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD'15), Porto, Portugal, pp 190–205

Neri M, Casu F, Acocella V, Solaro G, Pepe S, Berardino P, Sansosti E, Caltabiano T, Lundgren P, Lanari R (2009) Deformation and eruptions at Mount Etna (Italy): A lesson from 15 years of observations. Geophysical Research Letters 36(2):1–6

Nezry E, Genovese G, Solaas G, Rémondière S (1995) ERS - based early estimation of crop areas in Europe during winter 1994-95. In: Proceedings of the European Space Agency Second International Workshop on ERS Application, London, UK, vol 383, p 13

Nguyen T, Méger N, Rigotti C, Pothier C, Andreoli R (2016) SITS-P2miner: Pattern-Based Mining of Satellite Image Time Series. In: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, demonstration, (ECML-PKDD'16), Riva del Garda, Italy, pp 63–66

Novak PK, Lavrač N, Webb GI (2009) Supervised descriptive rule discovery: A unifying survey of contrast set, emerging pattern and subgroup mining. Journal of Machine Learning Research 10:377–403

Petitjean F, Masseglia F, Gancarski P, Forestier G (2011) Discovering Significant Evolution Patterns from Satelllite Image Time Series. International Journal of Neural Systems 21(6):15

Petitjean F, Inglada J, Gancarski P (2012) Satellite image time series analysis under time warping. IEEE Transactions on Geoscience and Remote Sensing 50(8):3081–3095

Pluim J, Maintz J, Viergever M (2003) Mutual-information-based registration of medical images: a survey. IEEE Transactions on Medical Imaging 22(8):986–1004

Proctor J (2003) Vegetation and soil and plant chemistry on ultramafic rocks in the tropical Far East. Perspectives in Plant Ecology, Evolution and Systematics 6(12):105 – 124

Quegan S, Toan TL, Yu JJ, Ribbes F, Floury N (2000) Multitemporal ERS-SAR analysis applied to forest mapping. IEEE Transactions on Geoscience and Remote Sensing 38(2):741–753

Rokni K, Ahmad A, Solaimani K, Hazini S (2015) A new approach for surface water change detection: Integration of pixel level image fusion and image classification techniques. International Journal of Applied Earth Observation and Geoinformation 34:226 – 234

Schellenberger T, Ventura B, Zebisch M, Notarnicola C (2012) Wet snow cover mapping algorithm based on multitemporal COSMO-SkyMed X-Band SAR images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 5(3):1045–1053

Sevin B, Maurizot P, Vendé-Leclerc M (2012) Geological map of New Caledonia Grand Sud. Map 1:50,000. Service Géologique de Nouvelle-Calédonie, Bureau de Recherches Géologiques et Minières (BRGM), Nouméa. `https://dimenc.gouv.nc/sites/default/files/download/grandsud.pdf` (accessed 08/25/2018)

SITS-Miner-team (2016) SITS-P2miner: a tool to build SITS summaries `https://sites.google.com/view/sits-p2miner` (accessed 08/25/2018)

Su X, Deledalle CA, Tupin F, Sun H (2014) Two-step multitemporal nonlocal means for synthetic aperture radar images. IEEE Transactions on Geoscience and Remote Sensing 52(10):6181–6196

Suri S, Reinartz P (2010) Mutual-information-based registration of terrasar-x and ikonos imagery in urban areas. IEEE Transactions on Geoscience and Remote Sensing 48(2):939–949

Tanimoto TT (1958) An elementary mathematical theory of classification and prediction. Tech. rep., Internal International Business Machines Corporation (IBM)

Tatti N, Vreeken J (2012) The long and the short of it: Summarising event sequences with serial episodes. In: Proceedings of the ACM 18th International Conference on Knowledge Discovery and Data Mining (KDD'12), Sydney, Australia, pp 462–470

Tedstone AJ, Nienow PW, Gourmelen N, Dehecq A, Goldberg D, Hanna E (2015) Decadal slowdown of a land-terminating sector of the Greenland Ice Sheet despite warming. Nature 526(7575):692–695

USGS, NASA, LANDSAT 7 Science Team (2003) Preliminary assessment of the value of Landsat 7 ETM+ data following Scan Line Corrector malfunction. `http://landsat.usgs.gov/documents/SLC_off_Scientific_Usability.pdf` (accessed 08/25/2018)

Vina A, Echavarria, R F, Rundquist (2004) Satellite change detection analysis of deforestation rates and patterns along the Colombia-Ecuador border. AMBIO: A Journal of the Human Environment 33:118–125

Viola P, Wells W (1995) Alignment by maximization of mutual information. In: Proceedings of the 5th International Conference on Computer Vision (ICCV'95), Coral Gables , Florida , USA, pp 16–23

Vreeken J, van Leeuwen M, Siebes A (2011) Krimp: mining itemsets that compress. Data Mining and Knowledge Discovery 23(1):169–214

Wang J, Han J (2004) BIDE: Efficient mining of frequent closed sequences. In: Proceedings of the 20th International Conference on Data Engineering (ICDE'04), Boston, MA, USA, pp 79–90

Wang Q, Jiang Yh, Zhang G, Sheng Qh (2015) Earthquake monitoring for multitemporal images of ziyuan-3. In: Proceedings of the SPIE International Conference on Intelligent Earth Observing and Applications, Guilin, China, vol 9808U, p 9

Wilcox M, Platt G (2002) Some observations on the flora of New Caledonia. Aukland Botanical Society Journal 57(1):60–75

Yan X, Han J, Afshar R (2003) Clospan: Mining closed sequential patterns in large databases. In: Proceedings of the Third SIAM International Conference on Data Mining (ICDM'03), San Fansisco, CA, USA, pp 166–177

Yu D, Wu W, Zheng S, Zhu Z (2012) BIDE-based parallel mining of frequent closed sequences with MapReduce. In: Proceedings of the International Conference on Algorithms and Architectures for Parallel Processing, Fukuoka, Japan, pp 177–186

Zhu X, Liu D (2014) Accurate mapping of forest types using dense seasonal landsat time-series. ISPRS Journal of Photogrammetry and Remote Sensing 96:1 – 11