Spatiotemporal Event Detection and Analysis over Multiple Granularities

Arie Croitoru, Kristin Eickhorst, Anthony Stefandis, Peggy Agouris

Department of Spatial Information Science & Engineering and National Center for Geographic Information and Analysis, The University of Maine, U.S.A; email: {arie, snoox, tony, peggy}@spatial.maine.edu

Abstract

Granularity in time and space has a fundamental role in our perception and understanding of various phenomena. Currently applied analysis methods are based on a single level of granularity that is user-driven, leaving the user with the difficult task of determining the level of spatiotemporal abstraction at which processing will take place. Without a priori knowledge about the nature of the phenomenon at hand this is often a difficult task that may have a substantial impact on the processing results. In light of this, this paper introduces a spatiotemporal data analysis and knowledge discovery framework, which is based on two primary components: the spatiotemporal helix and scale-space analysis. While the spatiotemporal helix offers the ability to model and summarize spatiotemporal data, the scale space analysis offers the ability to simultaneously process the data at multiple scales, thus allowing processing without a priori knowledge. In particular, this paper discusses how scale space representation and the derived deep structure can be used for the detection of events (and processes) in spatiotemporal data, and demonstrates the robustness of our framework in the presence of noise.

1 Introduction

Many phenomena in virtually all areas of natural sciences involve the study of change, and in particular change in spatial data over time. A primary reason for this interest in change is simple: change has a fundamental role in our perception and understanding of the world as it provides a systematic approach to the evolution of things in space and time. The identification and formalization of change patterns allows us to achieve what is often taken for granted: formalize rules, apply reasoning, and predict future behaviors of a given phenomenon. Consequently, the study of change in spatial data over time is essential in various areas, such as meteorology, geophysics, forestry, biology, and epidemiology.

The study of change in all these disciplines is closely related to the study of *events*. The description of change in terms of events is natural to us, primarily because as humans we intuitively tend to perceive an activity as consisting of discrete events (Zacks and Taversky 2001). Yet the term event, which is often used rather loosely in daily life, may have different meanings under different circumstances and different contexts. In view of their wide variability in space and time, there have been various suggestions for a more general model of events. For example, Galton (2000), based on an analysis of change, has identified three classes, namely states (the absence of change), processes (on-going change), and events (a predefined amount of change). Yet, as Galton has indicated, the distinction between states and events is not always straightforward: *"Thus processes seem to have a chameleon-like character, appearing now as states, now as events, depending on the context in which they are considered."* (Galton 2000, p 215).

Zacks and Taversky (2001) have also addressed the nature of events in human perception and conception and have defined events as a segment of time at a given location that is conceived by an observer to have a beginning and an end. Yet, in light of this definition they have also indicated that *"In general, it seems that as we increase the time-scale of our view, events become less physically characterized and more defined by the goals, plans, intentions and traits of their participants"* (Zacks and Taversky 2001, p 7).

These different examples of how events are defined and considered emphasize the intricate nature of events. In particular, a primary factor that contributes to the dual nature of events (processes⇔events and objects⇔events) is *granularity* in time and space, which has a fundamental role in our perception and understanding of various phenomena. Granularity refers to the notion that the world is perceived at different grain sizes

(Hornsby and Egenhofer 2002), and in the context of this work relates to the amount of detail necessary for a data analysis task.

The motivation for this work stems from the effect granularity has on our ability to analyze and understand spatiotemporal events. Currently applied analysis methods are based on a *single* level of granularity that is user-driven, that is, the user has the difficult role of determining a level of spatiotemporal abstraction at which the processing will take place. Yet in many cases, and in particular when analyzing unfamiliar phenomena, it is difficult to determine beforehand the granularity level at which event processing should take place without *a priori* knowledge: a too fine granularity will result in detail overloading while a too coarse granularity will result in over abstraction and loss of detail. In other cases, users may not be interested in events in a single granularity level, but rather in a range of levels in the study of events in given phenomenon (for example, users may be interested in analyzing rapid changes in the cloud mass of a hurricane both at the hourly and daily levels). This often leads to a need to repeat the processing of the data at each granularity level, thus resulting in low efficiency and high computational cost. Furthermore, such an approach makes it difficult to compare between phenomena that have a similar behavior but occur at different temporal scales.

It should be noted that the issue of the effect of granularity on the analysis of events is closely related to the problem of temporal zooming (Hornsby 2001) and the modeling of moving objects over multiple granularities (Hornsby and Egenhofer 2002). Yet, while this previous work focused primarily on the transition between identity states of objects and on adjusting the level-of-detail of object representations as a result of a change in granularity, the focus of this contribution is on the simultaneous analysis of spatiotemporal data over multiple granularities rather than a specific one.

In light of this, the main contribution of this work is the introduction of a novel spatiotemporal data analysis framework, which is based on two primary building blocks: *the spatiotemporal helix* – a formal representation of spatiotemporal data, and a *scale-space analysis* of the data in the temporal domain. Such an analysis offers two distinct advantages, namely, the ability to analyze the data at multiple granularities instantaneously, and the ability to reveal the hierarchical structure of events within the given phenomenon. To illustrate this, we will focus on a hurricane data analysis application, in which it is required to discover similar hurricanes by clustering. Our data source in this case is a time-series of remotely sensed imagery of each of the hurricanes, as provided by the National Oceanic and Atmospheric Administration (NOAA) (NASA 2005).

The remainder of this paper is organized as follows: Section 2 provides an overview of the spatiotemporal helix model that is used in this work to represent and summarize spatiotemporal phenomena. Section 3 provides a review of the scale space representation, followed by an analysis of the utilization of this representation as an event detection and analysis framework in spatiotemporal helixes. Section 4 describes an example application of the proposed framework using real-world data. Finally, conclusions and future work are summarized in Section 5.

2 The Spatiotemporal Helix

In our paradigm the spatiotemporal helix is used as a formal model of a spatiotemporal phenomenon. The spatiotemporal helix is a framework for describing and summarizing a spatiotemporal phenomenon. It is composed of a compact data structure and a set of summarizing tools capable of generalizing and summarizing spatiotemporal data, thus allowing efficient querying and delivering of such data. This framework can be applied to a variety of spatiotemporal data sources, ranging from manual monitoring of an object through time to spatiotemporal data that is collected through a single sensor or a sensor network (Venkataraman et al. 2004).

The idea behind the spatiotemporal helix is based on the observation that as an object moves, two key characteristics change over time: location and deformation. Using an image time series [see Fig. 1(a)] we extract the object using image-based feature extraction, from which we can track the object's location by calculating its center of mass and collecting this information in a database. In addition, we also track the deformation of the object by recording expansion and contraction magnitudes in each of four cardinal directions. This inclusion of deformation in the helix model provides the ability not only to track the changes in the location of the object, but also changes in its morphology. The result of the feature extraction process is depicted in Figure 1(b), which shows that a visualization of the extraction results in a three-dimensional space, consisting of a number of object out-lines stacked one on top of the next, with time as the vertical axis.

While initially all possible location and deformation information is gathered from the image time series, the summarization aspect of the helix is introduced by retaining only the frames that include significant information based on a user-defined change threshold. In this process we collect two types of entities: *nodes* consisting of the coordinates of the object's center of mass for each time instance, and *prongs* that capture information about the object's expansion or contraction. These *prongs* consist of a record of the time instance, magnitude, and direction in which the object has expanded or contracted. Nodes and prongs are therefore the building blocks of the spatiotemporal helix, where nodes construct the spine of the helix and prongs provide an annotation of the spine [see Fig. 1(c)]. An outline of the helix construction process from motion imagery is depicted in Figure 1. The interested reader may find more information about the spatiotemporal helix model in Agouris and Stefanidis (2003) and Stefanidis et al. (2005).

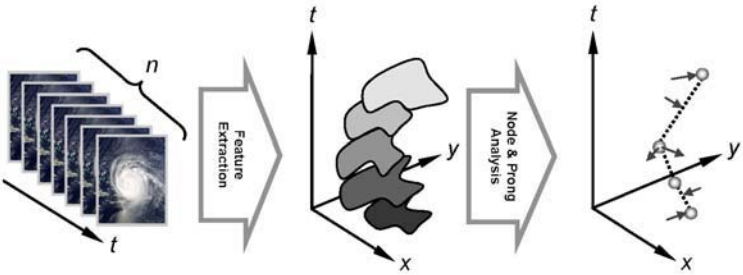

(a) (b) (c)

Fig. 1. Spatiotemporal helix visualized as stacking of objects over time. **(a)** An image time series of size *n*. **(b)** Feature extraction results in space and time. **(c)** The spatiotemporal helix (arrows represent prongs, circles represent nodes)

3 Scale Space Analysis

As was mentioned earlier, scale space analysis has a central role in our framework due to several distinct advantages it offers in the context of spatiotemporal event analysis. In order to demonstrate this we will first provide a short overview of scale space representation, and will then analyze the different characteristics in the context of spatiotemporal event analysis. It should be noted that in view of the ample body of literature on scale space representation, only a brief and non-exhaustive description of the key ideas and results are provided here. The interested reader may refer to Lindeberg (1994a, 1994b) and Sporring et al. (1997) for further details.

3.1 Scale Space Representation and Deep Structure

The development of the scale space representation stemmed from the understanding that scale plays a fundamental and crucial role in the analysis of measurements (signals) of physical phenomena. In order to demonstrate this, consider a signal *f*, which was obtained from a set of real-world measurements. The extraction of information from *f* is based on the application of an operator with a predefined scale. A fundamental question in this process is therefore the determination of the *proper scale* of the operator. Clearly, there is a direct connection between the scale of the operator we chose to apply and the scale of the structures (information) in *f* that we wish to detect (Lindeberg 1994b). If the scale of the operator is too large or too small, our ability to derive information from *f* will be compromised, leading to either high sensitivity to noise or low sensitivity to the structures sought. Consequently, proper scale should be used in order to ensure the optimal extraction of meaningful information from *f*. The determination of the proper scale is straightforward in cases where a priori knowledge about *f* exists, yet in other cases where there is no a priori knowledge the determination of the proper scale becomes a fundamental challenge and all scales should be considered. This notion of considering all possible scales is at the heart of the scale space representation (Lindeberg 1994b).

The construction of a scale space representation is carried out by embedding the signal *f* into a one-parameter family of derived signals, in which the scale is controlled by a scale parameter σ. More formally, given a signal $f(x):\Re\rightarrow\Re\ \forall x\in\Re$, the (linear) *scale space representation* $L(x,\sigma):\Re\times\Re+\rightarrow\Re$ of $f(x)$ is defined such that $L(x,0)= f(x)$, and the representation at coarser scales are given by $L(x,\sigma)=g(x,\sigma)* f(x)$, where $*$ is the convolution operator, and $g(x,\sigma)$ is a Gaussian kernels of increasing width (Lindeberg 1994b). In the case of a one-dimensional signal, $g(x,\sigma)$ is taken as the one-dimensional Gaussian kernel (Witkin 1983; Lindeberg 1990):

$$g(x,\sigma)=\frac{1}{\sqrt{2\pi\sigma}}e^{-x^2/2\sigma} \tag{1}$$

The two-dimensional space formed by (x,σ) is termed *scale space*. The scale space representation of $f(x)$ is therefore comprised from a family of curves in scale space that have been successively smoothed by the kernel.

While the generation of the scale space representation (L) results in a family of signals with an increasing level of smoothing, it is the inner structure of the scale space representation that exhibits distinct inherent behavior. In particular, it was found that the extrema points (the zero-crossings of the n^{th} derivative) in scale space representation form paths in scale space that will not be closed from below and that no new paths will be created as σ increases. Hence, as σ increases new extrema points cannot be created (Witkin 1983; Mokhtarian and Mackworth 1986). We carry out the generation of such paths by computing the location of the zero-crossings for each of the derived signals in L, and then stacking these different locations in scale space. As result of this process a binary image showing these paths is created. Following Florack and Kuijper (2000) and Kuijper et al. (2001), we term the resulting binary image the *deep structure* of the Gaussian scale space, that is the structure of at all levels of granularity simultaneously.

To illustrate how the scale space representation and the deep structure are used for the detection of features in the data consider the example depicted in Figure 2, which shows the analysis of a one-dimensional signal [see Fig. 2(a) top].

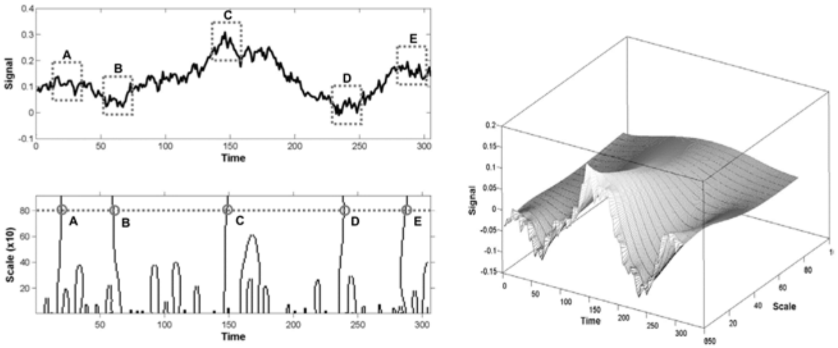

(a) (b)

Fig. 2. An example of a scale space analysis of a one-dimensional signal. **(a)** Top – A one-dimensional signal; Bottom – the deep structure of the signal as de-rived by the zero-crossings of the 1st derivative **(b)** The scale space representation of the signal, which was derived using a Gaussian kernel of increasing size. In all figures the *x*-axis represents time

By convolving this signal with a Gaussian kernel of an in-creasing size (see Eq. 1) the scale space representation (L) is derived. Figure 2(b) shows the scale space representation as a three-dimensional surface (note how the original signal becomes smoother as the scale of the Gaussian kernel increases). Once the scale space representation is derived, the zero-crossings of the n^{th} derivative are detected for each scale level and the deep structure [see Fig. 2(a), bottom] is recovered by stacking the zero-crossings one on top of the other. Here, for the sake of simplicity, we have chosen to use the zero crossing of the 1st derivative to construct the deep structure due to its direct relation to features in the signal. Consequently, the paths formed in the deep structure describe how extrema points in the signal evolve as scale increases. To demonstrate this, consider the deep structure at a scale level of 8 [marked by the dashed horizontal line in Fig. 2(a), bottom] which shows five points, A through E. Clearly these five points correspond to local extrema in the signal, as marked by the five rectangles in Figure 2(a), top. It is easy to see that as the scale parameter decreases more extrema points appear in the deep structure due to the noisy nature of the signal.

3.2 Analysis of Events Using Scale Space Representation

The adaptation of the scale space representation approach to the analysis of events in spatiotemporal helixes can offer several distinct advantages. In order to demonstrate this we first define events within the helix framework and then analyze the different characteristics of the scale space representation in light of this definition.

Following Galton (2000) and Grenon and Smith (2004), we define events within the context of spatiotemporal helixes as the *transition between states*. As such, we regard events as all entities that exhaust themselves in a single instant of time (Grenon and Smith 2004). Consequently, events are used to define the boundaries of processes and indicate the transition within processes.

Having this definition in mind, let us now analyze how the scale space representation and the deep structure of a physical measurement signal could be used for the detection of events. As noted earlier, the zero crossing of the n^{th} (commonly n=2) derivative of L is used to construct the deep structure. Since the zero-crossing condition ensures a change of sign in the second derivative (note that this is not the same as requiring that the second derivative will be zero), the deep structure serves as an indicator of inflection points[1] in the given signal f. Such inflection points indicate either a minimum or a maximum in the gradient of L. In conclusion, the deep structure can be used for detecting minimum or maximum rates of change (the first derivative) of a process, or changes in the sign of the rate of change (the second derivative) of a process. Note that in our interpretation we view processes as occurring between inflection points (events), which corresponds to our definition of events.

To illustrate this, let us assume that f is a vector of the x coordinate of a hurricane that was tracked in time. The first derivative of f indicates the speed of the hurricane in the x direction, while the second derivative indicates the acceleration of the hurricane in the x direction. Transforming f into a scale space representation and the recovery of the deep structure en-

[1] Given a twice differentiable function $g(x)$, a point $x=c$ on g is an inflection point if $g(c)'$ is an extremum point and $g(c)''$ changes its sign in the neighborhood of c (Binmore 2001).

ables the detection of the following events: *maximum speed events, minimum speed events, acceleration to deceleration events, and deceleration to acceleration events.*

Let us now turn back to the deep structure of spatiotemporal helixes and analyze its characteristics in relation to the analysis of events. Several key observations can be made here:

1. As was mentioned earlier, the derived deep structure is based on the detection of inflection points, which form paths that are guaranteed not be closed from below. It is also guaranteed that no new paths will be created as σ (the scale factor) increases. These characteristics are essential in the analysis of events as it is expected that no new processes will emerge, as the time granularity of a physical process is made coarser. Furthermore, this property ensures that a process (which is defined between inflection points) cannot disappear and then reemerge in coarser time granularities.
2. In general, paths in the deep structure will not cross each other (Mokhtarian and Mackworth 1986). This property of the paths assures that time conflicts will not occur. Consider for instance the example in Figure 2: since the x axis of both the data and the deep structure is time, paths that cross each other will indicate that two events that occurred in one order in one time granularity level will occur in the opposite order in another granularity level. This property therefore ensures that the proper order of events will be maintained at all granularity levels.
3. In general, paths in the deep structure will either form arch-like paths that are closed from above [for example, curve e in Fig. 3(a)], or will form a single path line. In the context of events, these properties indicate that as time granularity is made coarser processes that are defined between two events (inflection points) will converge to a single *event point* (the top point of a path for which the gradient is zero) and eventually disappear. This can be seen in Figure 3(a), where path b converges to a single point (tangent to the dashed horizontal line) at scale σ_i, and disappears at coarser granularities. In addition, single path lines indicate *transition events* [for example, curve a in Fig. 3(a)] that are not defining processes within the framework of the data provided but rather a change in the process.

In summary, the deep structure (and the scale space representation) can be used for the detection of events through which processes can be defined. Furthermore, the Gaussian scale space ensures that as time granularity is made coarser (a) no new processes will emerge, (b) processes can not disappear and reemerge, (c) the proper order of events (and processes) is maintained, (d) processes in lower granularity will tend to converge to a single event point, and eventually disappear.

In addition to these characteristics it is important to note that the deep structure inherently offers the ability to reveal the *hierarchy* of events and processes. To illustrate this, consider the scale space path b in Figure 3(a). As can be seen, this path contains two additional paths, c and d, which can be seen as two *sub-processes*. It should be noted that as granularity increases sub-processes turn to a single point event and eventually disappear. This hierarchy can be further described in a process tree [see Fig. 3(b)].

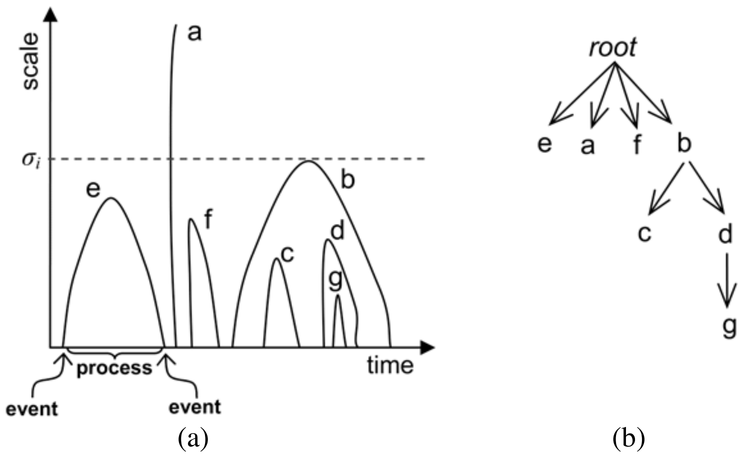

Fig. 3. Event and process hierarchy discovery through scale space representation. **(a)** A sample deep structure. **(b)** The derived process tree

3.3 Scale Space Analysis of Events in Spatiotemporal Helixes

As was described in Section 2, the spatiotemporal helix is a framework for describing and summarizing a spatiotemporal phenomenon. Given an image time series which contains an object that should be analyzed, the spatiotemporal helix collects the following information about the object along its spine (Stefanidis et al. 2005): the *x* and *y* location of the center of mass, acceleration, rotation, and expansion/contraction in north, east, south, and west directions.

In order to analyze the helix and detect events (and processes), we treat each of the 8 attributes of the spatiotemporal helix as a one-dimensional signal. Based on this, we can then apply a scale space analysis for each of these signals. This will result in 8 scale space representations and deep structures that could be used in various ways in order to understand the underlying physical phenomenon at hand. In particular, the following applications of the deep structures are considered in the analysis of helixes:

1. *Single helix dimension analysis* – in this case, a single dimension of a single helix can be analyzed (for example, the x or y location of the center of mass). Here, the deep structure can be used as an inspection tool for the detection of events and processes at multiple granularities. Furthermore, the deep structure provides an insight into the evolution of events and to the *hierarchy* of processes [see Fig. 3(b)].
2. *Multiple helix dimension analysis* – in this case, two or more dimensions of a single helix are analyzed by overlaying their deep structures. This would allow, for instance, detecting of processes that occur at the same time interval in different dimensions, from which higher-level inferences about the phenomenon could be derived.
3. *Helix clustering* – in this case, the primary goal is to estimate the similarity between helixes for the purpose of discovering similar physical phenomenon. In this case the similarity function, $S(\cdot,\cdot)$, between helixes H_i and H_j can be computed by:

$$S\left(H_i, H_j\right) = \sum_{k=1}^{n} w_k C\left(DS_k\left(H_i\right), DS_k\left(H_j\right)\right) \tag{2}$$

where k is the number of dimensions (attributes) in each helix, w_k is a weight assigned to each dimension (user defined), $DS_k(\cdot)$ is the deep structure of the k_{th} dimension (a two-dimensional matrix of size $u\times v$), and $C(\cdot,\cdot)$ is the two-dimensional cross correlation coefficient between $DS_k(H_i)$ and $DS_k(H_j)$ that is given by:

$$C\left(DS_k\left(H_i\right), DS_k\left(H_j\right)\right) = \frac{\sum_u \sum_v \left(DS_k\left(H_i\right) - \overline{DS}_k\left(H_i\right)\right)\left(DS_k\left(H_j\right) - \overline{DS}_k\left(H_j\right)\right)}{\sqrt{\sum_u \sum_v \left(DS_k\left(H_i\right) - \overline{DS}_k\left(H_i\right)\right)^2}\sqrt{\sum_u \sum_v \left(DS_k\left(H_j\right) - \overline{DS}_k\left(H_j\right)\right)^2}} \tag{3}$$

If the length of the two helixes is not the same, a template matching approach is used, in which the deep structure of the shorter helix is shifted along the x axis (time) of the deep structure of the longer helix, and the maximum cross correlation coefficient [see Eq. (3)] is taken. It should be emphasized that because the deep structure of the k^{th} dimension is used in Equation (3), the cross correlation is being simultaneously computed in multiple granularities. Thus, S [see Eq. (2)] is a measure of the overall similarity of the two helixes over multiple granularities.

4 An Example: Spatiotemporal Hurricane Data Clustering

In order to demonstrate the capabilities and robustness of our approach we have applied the proposed framework to the following problem: *given a data set of n imagery time series of physical phenomena, partition the data set into subsets of similar phenomena without any a priori knowledge about the granularity of the phenomena*. The primary motivation for selecting this particular task was the centrality that role clustering has in numerous areas, such as data mining and knowledge discovery, search engines, machine learning, and pattern recognition. In all these areas reliance on minimal a priori knowledge and robustness to noise are crucial.

For this work we have collected real-world satellite imagery time series of five different tropical storms and hurricanes. The satellite data was obtained from NASA-GFSC's GOES project (http://goes.gsfc.nasa.gov) that provides GOES-12 imagery. The storms that were collected are Alex (1–5 August, 2004), Allison (4–14 June, 2001), Charley (11–14 August, 2004), Dennis (7–11 July, 2005), and Frances (August 31 – September 7, 2004).

The preliminary processing of each of the five hurricane image time series included the delineation of the contour of the hurricane cloud mass from each image frame. This process, which resulted in a binary image time series, was then used as input to the helix construction process (Stefanidis et al. 2005) from which a spatiotemporal helix was created for each hurricane (see Fig. 4). Then, from each of the five helixes four more permutations were created by corrupting the original helix data with an increasing level of random noise that was added to the center of mass and the expansion/contraction dimensions. This process resulted in a data set consisting of a total of 25 helixes. In order to cluster this data set we have implemented two different methods:

1. *The local extreme event approach* – in this approach we defined extreme events based on the deviation from the average attribute value using a moving window, that is, given a confidence level d all helix attribute values within a window of a user-defined size that deviate more than $d \cdot \sigma$ from the average are considered to be extreme events. By changing the window size and the value of d the user can then control the level of granularity in which extreme events are detected. Based on these extreme events we then computed the distance between all possible helix pairs using the technique described by Stefanidis et al. (2005), and constructed a distance matrix from which a dendrogram was derived.
2. *The scale space approach* – in this approach we implemented the pro-posed scale space clustering technique that was described in Sec-

tion 3.3. Similar to the first approach, here we have also computed the distance between all possible helix pairs using Equation (2) and (3), and constructed a distance matrix from which a dendrogram was derived.

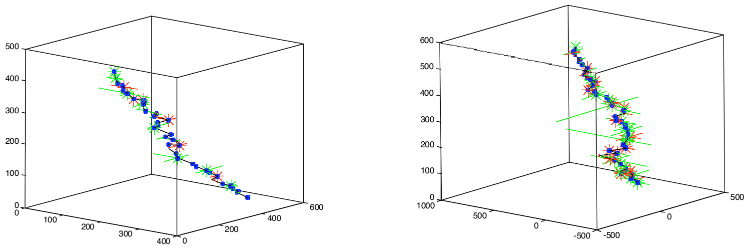

(a) (b)

Fig. 4. Examples of the derived spatiotemporal helixes for the hurricane data sets. **(a)** hurricane Frances, **(b)** hurricane Alison. In both figures the central black line represents the spine of the helix, the black circles represent the nodes, and the gray lines represent prong information

Using these methods two different clustering experiments were conducted. The first experiment included the clustering of data from two hurricanes, Frances and Alison, including their permutations (a total of 10 helixes) using both methods. The second included the clustering of the entire hurricane data set (a total of 25 helixes) using the scale space approach. In both experiments a correct clustering would result in distinct clusters in the dendrogram, where each cluster contains data from only one of the hurricanes. An example of the deep structure (2^{nd} derivative) that was derived and utilized for each helix in both experiments is depicted in Figure 5.

The results of the first experiment are depicted in Figure 6, where Figure 6(a) through (c) show the distance matrices and dendrograms that were obtained from the local extreme event approach.

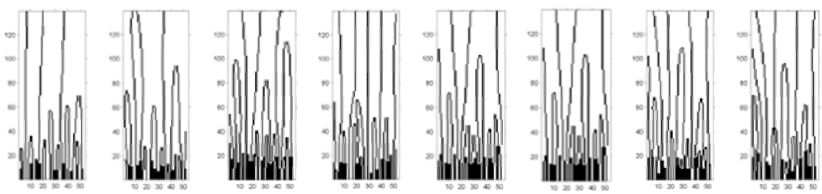

(a) (b) (c) (d) (e) (f) (g) (h)

Fig. 5. The deep structure of the spatiotemporal helix of hurricane Alison. **(a)** *x* location, **(b)** y location, **(c)** acceleration, **(d)** rotation, **(e)** through **(h)** expansion / contraction in the north, east, south, and west direction respectively

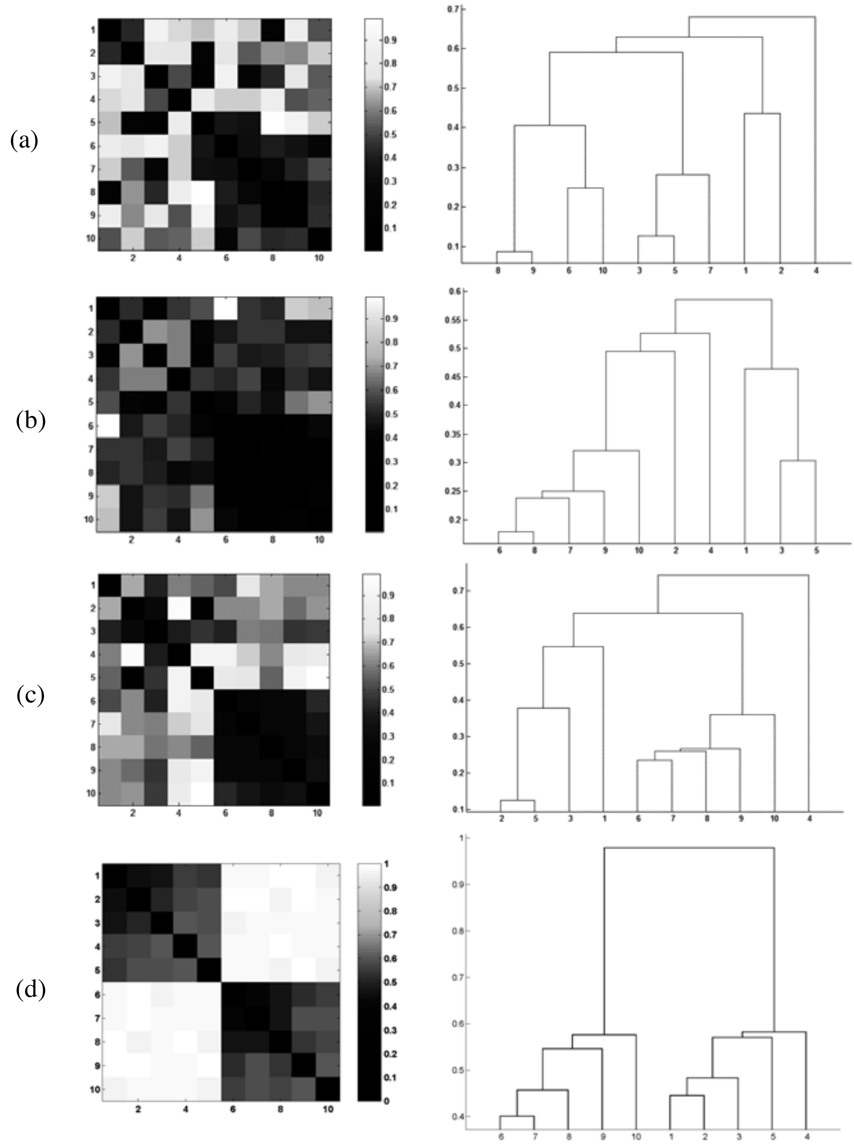

Fig. 6. Results of the first clustering experiment. **(a)** through **(c)** – The distance matrix (left) and dendrogram (right) using the local extreme event approach with a moving window size of 5, 7, and 9 respectively. **(d)** The distance matrix (left) and dendrogram (right) using the scale space approach. In all figures numbers 1–5 correspond to Allison and numbers 6–10 correspond to Frances. Darker shades in the distance matrices correspond to higher similarity

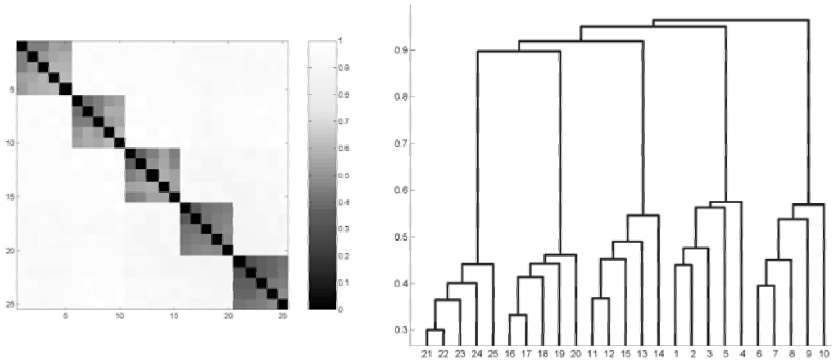

(a) (b)

Fig. 2. Results of the second clustering experiment using the entire hurricane data set. **(a)** The distance matrix. **(b)** The resulting dendrogram. In both figures numbers 1–5 correspond to Allison, 6–10 correspond to Frances, 11–15 correspond to Dennis, 16–20 correspond to Alex, and 21–25 correspond to Charley

The effect of the granularity at which the processing takes place is evident: as the window size increases some clusters do begin to emerge; yet the correct clustering is not obtained. In practical application this demonstrates the difficulty users are likely to face when analyzing such data without proper a priori knowledge. In contrast, the scale space approach produced the correct clustering [see Fig. 6(d)], resulting in two well-defined clusters, one for each set of hurricane data.

The results of the second experiment are depicted in Figure 7. As can be seen, the scale space approach successfully recovered the five clusters of hurricanes in this case as well.

5 Conclusions

Granularity in time and space has a fundamental role in our perception and understanding of various phenomena. Furthermore, since improper granularity may lead to erroneous results, it is essential that proper granularity be used in spatiotemporal data analysis and knowledge discovery. In spite of the importance of granularity, it is often difficult to determine at which granularity data processing should take place without a priori knowledge. This paper addressed this problem by adopting a scale space approach, in which all granularity levels are considered instead of applying a single granularity level. Based on this approach we presented a framework consisting of the spatiotemporal helix as a modeling and summarization tool, and the scale space representation as an analysis and knowledge discovery tool. The primary advantage of our framework is that it does not require a priori knowledge about granularity.

We analyzed how scale space representation and the derived deep structure could be used for the detection and analysis of events and processes and showed that due to its unique characteristics, deep structure can be used for the detection of events through which processes can be defined. Furthermore, we showed that the deep structure ensures that as time granularity is made coarser no new processes will emerge, processes can not disappear and reemerge, the proper order of events (and processes) is maintained, and that processes in lower granularity will tend to converge to a single event point, and eventually disappear. Additionally, we described how the deep structure could be used for the discovery of a hierarchy of events and processes. To demonstrate the capabilities of our approach we applied the proposed framework to the problem of real-world hurricane data clustering, and showed its robustness in the presence of noise.

In the future we plan to further explore and expand our framework. In particular, we are interested in utilizing the proposed approach for determining the proper granularity that should be used in the analysis of a given data set, and in developing additional similarity functions for scale space representations.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the National Geospatial-Intelligence Agency through grant NMA 401-02-1-2008, and by the National Science Foundation through grant ITR-0121269.

References

Agouris P, Stefanidis A (2003) Efficient summarization of spatiotemporal events. Communications of the ACM 46(1):65–66

Binmore KG (2001) Calculus. Cambridge University Press

Florack L, Kuijper A (2000) The topological structure of scale-space images. J of Mathematical Imaging and Vision 12:65–79

Galton A (2000) Qualitative spatial change. Oxford University Press

Grenon P, Smith B (2004) SNAP and SPAN: towards dynamic spatial ontology. Spatial Cognition and Computation 4(1):69–104

Hornsby K (2001) Temporal zooming. Transactions in GIS 5(3):255–272

Hornsby K, Egenhofer M (2002) Modeling moving objects over multiple granularities, Special issue on Spatial and Temporal Granularity. In: Annals of Mathematics and Artificial Intelligence 36. Kluwer Academic Press, Dordrecht, pp 177–194

Kuijper A, Florack LMJ, Veirgever MA (2001) Scale space hierarchy. Technical report UU-CS-2001-19, Utrecht University, Department of Computer Science

Lindeberg T (1990) Scale space for discrete signals. IEEE Transactions on Pattern Analysis and Machine Intelligence 12(3):234–254

Lindeberg T (1994a) Scale-space Theory: A Basic Tool for Analyzing Structures at Different Scales. J of Applied Statistics 21(2):225–270

Lindeberg T (1994b) Scale space theory in computer vision. Kluwer Academic Press, Dordrecht

Mokhtrian F, Macworth A (1986) Scale based description and recognition of planar curves and two-dimensional shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 8(1):34–43

NASA (2005) The GOES science project. http://rsd.gsfc.nasa.gov/goes/goesproject.html (last visited: December 2005)

Sporring J, Florack L, Nielsen M, Johnsen P (1997) Gaussian scale-space theory. Kluwer Academic Publishers

Stefanidis A, Agouris P, Eickhorst K, Croitoru A (2005) Modeling Object Movements and Changes with Spatiotemporal Helixes. Submitted to the Int J of Geographical Information Science

Venkataraman V, Srinivasan S, Stefanidis A (2004) Object Color Propagation in an Unregistered Distributed Video Sensor Network. IEEE Int Conf on Image Processing (ICIP) 2004, Oct 2004, Singapore

Witkin AP (1983) Scale-space filtering. Proc of the 8th Int Joint Conf on Artificial Intelligence, Karlsruhe, Germany, pp 1019–1022

Zacks JM, Tversky B (2001) Event structure in perception and conception. Psychological Bulletin 127(1):3–21