Capturing and Modeling Geographic Object Change: A SpatioTemporal Gazetteer Framework

Peggy Agouris, Kate Beard, Georgios Mountrakis, and Anthony Stefanidis

Abstract

A SpatioTemporal Gazetteer *as a model to manage change information at the geographic entity instance level is presented. The Gazetteer framework links an image repository, that is the source of change information, to instances of geographic entities, and is change specific to these instances. The ability to quickly detect and extract change from image data is an essential component of the model. The paper also presents a modified template matching function used to detect change, and describes the association of this function with the spatiotemporal gazetteer. The change detection function embedded within the spatiotemporal gazetteer creates a foundation for up-to-date geospatial information and change information that can serve rapid-response decision-making applications.*

Introduction

Change, whether it is the very dynamic change of mobile objects, periodic fluctuations of fixed objects, abrupt change such as that resulting from catastrophic events, or more routine incremental change, is valuable information to which an information system should provide direct access. Information on change is important to better understand why certain conditions exist — either now or in the past — and to make predictions on what the conditions or configurations of an entity may be at some future time. Tools to model, analyze, and predict various temporal behaviors have broad applicability in geospatial analysis.

Here we present a *SpatioTemporal Gazetteer* (STG) as a model that makes use of multiple information resources and, in particular, incorporates components to track changes to objects over time. This Gazetteer framework manages representations of instances of geographic entities and their changes over time rather than changes to layers or scenes. It allows us to organize geospatial information in an object-oriented manner that captures essential components of the spatiotemporal behavior of objects. This approach matches the approach emphasized in the National Imagery and Mapping Agency's (NIMA's) recently announced vision for Integrated Information Libraries (NIMA, 2000). This vision is characterized, among other things, by

- the need to arrange information content as objects instead of thematic layers,
- the assignment of an ID to each object (manmade or natural) contained in the information environment, and
- the ability to arrange information content by dimensions in a variety of ways to better respond to specific user needs.

In this paper we present an overview of the *SpatioTemporal Gazetteer* and a modified template matching approach used to capture changes to geospatial objects. The paper is organized in the following manner. The next section discusses the role of our *SpatioTemporal Gazetteer* framework in support of decision making in applications requiring rapid access to changes in geospatial information. This is followed by a presentation on the *SpatioTemporal Gazetteer* framework, its components, and the way they are interrelated. The next section presents the image matching tool used to capture change and store it within the *Gazetteer* environment. Experimental results and final conclusions are then presented.

Decision-Making Support

A model for capturing and managing change information at a geographic entity instance level can be an asset for rapid response decision making where quick access to change information can be critical. Change information is useful to a wide range of application areas. Applications extend from global change studies involving subtle spatial changes over long time periods to rapidly occurring changes in an emergency response or military situation. Decisions that revolve around change must consider if, in fact, a change has occurred, what type of change occurred, if this change is significant, and its implications. In this paper we focus specifically on geometric change and a procedure for detecting such change.

Many decision-making environments require information on change, but current GIS do not deliver change information directly. The questions from a team monitoring an oil spill might be: Has the spill moved? Has it changed direction? Has it spread, split, or become fragmented? Another context for change-detection and analysis support involves military maneuvers. In an active battlefield environment many changes can be occurring that affect decisions on where to place and move troops. This change may be occurring over a wide range of resolution. Coarse geometric change such as loss or gain of whole structures may be easily obtained, but more subtle geometric changes can be equally important. Small changes in bridge geometry, for example, can have major impact on troop movements. Troops in the field need current and accurate

P. Agouris, K. Beard, and G. Mountrakis are with the Department of Spatial Information Engineering, and A. Stefanidis is with the National Center for Geographic Information and Analysis (NCGIA), all at 5711 Boardman Hall #348, University of Maine, Orono, ME 04469-5711 ({peggy; beard; giorgos; tony}@spatial.maine.edu).

Photogrammetric Engineering & Remote Sensing
Vol. 66, No. 10, October 2000, pp. 1241–1250.

0099-1112/00/6610–1241$3.00/0

information on bridge bed width, as well as other attributes, to quickly determine possible deployment routes.

In each of these cases, questions revolve around specific instances of objects and whether these objects have changed. GIS currently focus on queries of the type: What is here? and Where are things like this? While they support an array of analytical functions, without substantial customization they do not support queries on change such as: Has the wetland increased in size since last year? Direct access to change information on instances of geographic entities requires more than differencing images and reporting such differences. This paper describes an approach for detecting and managing geometric change information at the instance level, and at multiple levels of detail or resolution such that information may be readily available for decision makers in rapidly evolving field environments.

In our approach, change refers to types of changes that affect instances of real world geographic entities rather than changes that affect information sources such as atmospheric effects, time of day, varying resolutions, and other variations in the data collection mechanism. Change to geographic objects is categorized along three independent dimensions: movement, boundary reconfiguration, and thematic state change. The first two are considered geometric changes, and are the main focus of this paper. Either the spatial position of the object has changed (movement has occurred) or something in its spatial definition (the boundary that separates an instance from its surroundings) has changed. The third category refers to change in non-spatial attributes such as name, color, or weight of an object.

Our approach depends on rapidly detecting change using new imagery and storing what we call change primitives. As an operational environment, we could imagine a battlefield commander planning a move to a new location. We assume some initial spatial information is available for the area, and that new imagery and other surveillance information are being routinely acquired. The commander could request changes to specific geographic objects or classes of objects, request a specific type of change (has anything moved?), or ask for change in general (what if anything has changed in the area in the last 24 hours?). The goal is to provide responses that return change information at the level requested or, in the case of a general change request, to provide a response with a meaningful level of detail given the query semantics. A response to the above general query (what if anything has changed in the area in the last 24 hours?) could be that three bridges are destroyed and the water level in the river is up two feet. A follow-up question by the commander might be; Do these changes affect the ability to get troops from A to B? This scenario demonstrates the utility of having fast access to change information. Multiple sources of information collected and stored over short intervals contribute to the realization of this scenario, and many of these sources are imagery. In this paper we begin by describing the overall change model and focus specifically on the exploitation of imagery to detect geometric change.

The SpatioTemporal Gazetteer Framework

In order to model change to instances of geographic entities, we use a spatiotemporal gazetteer. The gazetteer stores representations of real-world geographic objects and representations of changes to these objects. Our spatiotemporal model is roughly based on a digital library model (Beard *et al.*, 1997; Beard and Smith, 1998; Goodchild, 1998). The approach shares some of the characteristics of the multimedia geographic information system proposed by Lombardo and Kemp (1997) and symbolic description of image sequences using spatiotemporal logic (Del Bimbo *et al.*, 1995). Before describing the structure of the gazetteer, we present our conceptual model of change that forms the basis of the gazetteer.

Conceptual Model of Geographic Entity Changes

The entity relation diagram shown in Figure 1 indicates the primary entities and relations among them that are used to model change information. The primary entities are

- spatial information sources which can include maps, images, GPS observations, video, etc.; and
- geographic entities which are real world objects such as bridges, roads, streams, buildings, vehicles, people, animals, etc.

Geographic entities come into being and may subsequently expire, be razed, or otherwise go out of existence. Existence is a fundamental change (Hornsby, 1999) indicated by a geographic entity's life-span (the time interval between an entity's creation and termination). During its life-span, an entity may experience other changes that can be categorized along a variety of dimensions. Considering standard GIS practice, changes in boundary, position, and thematic information are of high importance. Information on these three dimensions of change exists inherently in sets of spatial information sources collected over time.

Each information source represents an observation on one or more geographic entities for a point or interval of time. Collectively, the information sources contain multiple representations for a single geographic entity. These multiple representations of an entity may be due to variations in characteristics of the information sources (different scales, different resolutions, different observation perspectives, different formats (e.g., text versus graphic)) or due to change in the geographic entity with the passage of time. These one-to-many relations are captured in the **represents** and **is represented by** relations between a geographic entity and an information source as shown in Figure 1. For example, a single geographic entity may be observed in numerous aerial photos and maps, and within a single satellite image we may observe numerous geographic entities.

As illustrated on the diagram, we assume each geographic entity has one and only one boundary at a given time. This boundary, however, may be observed and recorded differently by different information sources, resulting in multiple boundary representations. Similarly, a geographic entity has one and only one position that it occupies at any given time, but this position may have multiple representations as given by different information sources. A geographic entity may have more

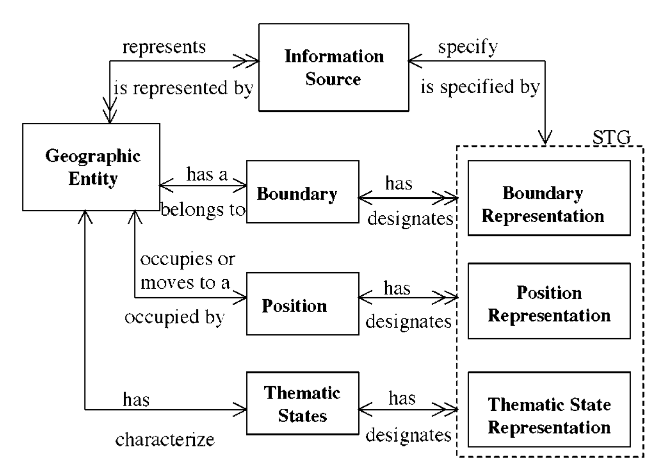

Figure 1. Entity relation diagram of geographic entity change. Rectangles represent entities. Connections represent relations in both directions between entities. Single-headed arrows represent one-to-one relations. Double-headed arrows represent many-to-one relations.

than one thematic state (e.g., name, color, or weight) but only one value for a given thematic state at a given time. Each thematic state may have multiple representations as given by different information sources.

The entity relation diagram of Figure 1 represents a single time instance. Over time an entity can modify one or more of its properties. It can modify its boundary, move to and occupy a new position, and transform to new thematic states. Over an interval of time, then, an entity may have multiple boundary configurations, multiple positions, and multiple states that are the result of change rather than variations in the information sources. The *SpatioTemporal Gazetteer* explicitly stores these multiple representations and employs methods to detect differences in representations that are due to change. As a result, change information is accessible at the individual feature level and separable into categories of position change, boundary change, and thematic state change.

Components of the SpatioTemporal Gazetteer Framework

The *SpatioTemporal Gazetteer* framework is an implementation of the entity relation model presented earlier. A gazetteer is commonly identified as a mapping from geographic feature names (place names) to spatial locations (footprints) and vice versa (Hill *et al.*, 1999). Our spatiotemporal gazetteer more broadly includes the mapping from instances of geographic entities (each uniquely identified) to multiple boundary representations, multiple location representations, and multiple thematic state representations. In each case the boundary, position, and thematic state representations are extracts from information sources. The gazetteer thus provides a mapping from information sources that may be stored as layers, themes, and scenes into a structure organized by individual geographic entities. Inherently included in this mechanism is a mapping from any geographic entity to its multiple representations.

Our model comprises a repository of heterogeneous information sources, with a set of indexing structures to organize and access them. The gazetteer is organized by what we call registers, with one register for geographic entities, one for their boundary representations, one for their position representations, and one for their thematic state representations (see Figure 2).

Tightly associated with the Gazetteer is what we call the Multimedia Information Store, a database of information sources. The Multimedia Information Store (MIS) includes imagery, maps, video, various types of scientific data sets (e.g., samples from field observations), and even digital files of text documents (e.g., books, newspaper reports, and magazines). The multimedia information store is a continually growing repository, with new information added on an ongoing basis.

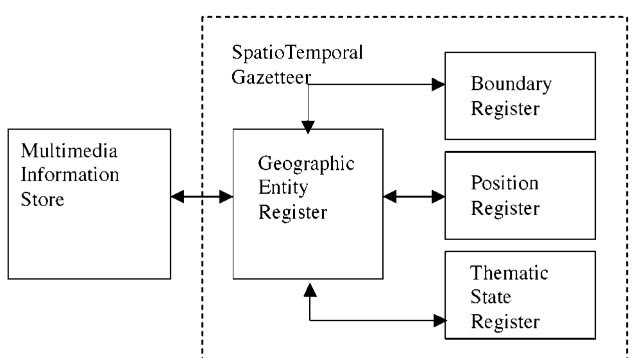

Figure 2. The SpatioTemporal Gazetteer framework.

This repository need not be a single site, but could be a multinode distributed repository. The National Spatial Database Infrastructure (NSDI) of the Federal Geographic Data Committee (FGDC) is an example of such a virtual repository. Each information source in the multimedia information store is represented as an information object with a minimum set of attributes that includes a unique identifier, an information source class, a timestamp, and a spatial footprint as shown in Figure 3.

The information source class specifies the type of information object (e.g., map, satellite image, aerial photograph, video, etc.) and is the basis for assigning class-specific metadata (e.g., resolution for satellite imagery, flying height for an aerial photograph). It should be noted here that, in our modeling, a timestamp is associated with each information object from the Multimedia Information Store, and is propagated to geographic objects within the *SpatioTemporal Gazetteer.* This timestamp represents the valid time of an event (Salzberg and Tsotras, 1999), which in our case reflects the time that the information object was captured. The time stamp for an information object can be a point or interval reflecting the time over which the information was collected. Some information sources (sensors) are designed to record a change of state the moment it occurs (e.g., a flood gauge), in which case the assigned time stamp corresponds to the instant of change. Other information sources can document changes of state but do not capture the exact moment that change occurred. This is typically the case with satellite imagery. As an example, an image can record that a house burned down but not necessarily when. Attributes of information objects can be used to indicate whether a specific source is capable of recording the exact time of a state change.

The spatial footprint represents the geographic area to which an information source refers, and it reflects the degree to which the geographic location is known. For spatial information objects such as maps, aerial photographs, and satellite images, the footprint is specified by the bounding coordinates of the map, scene, or image. Less precise footprints may still be generated for information objects that only reference geographic regions. For example, a research paper describing fisheries in Downeast Maine could be associated with a broad spatial footprint covering the general region.

While the information objects represent the information sources, the gazetteer components represent individual geographic entities and their depictions in information sources. The four register components of the gazetteer are described in more detail in the following sections.

Geographic Entity Register

The central component of the SpatioTemporal Gazetteer is the Geographic Entity Register. It is a database of all geographic entities of interest which have been identified as unique instances of features in the landscape. At the same time, it provides links between objects and datasets in which they are included. Each geographic entity is assigned a unique identifier to distinguish it from all other geographic entities. While

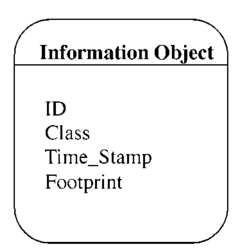

Figure 3. Minimum attribute set for an information object.

mechanisms to assign ID numbers are beyond the scope of this paper, one could easily envision a variety of approaches to do so (ranging from simple sequential numbering to complex strategies that combine sequential numbering with object type and geographic region coding). Beyond the identifier, the other attributes of a geographic entity include a life-span (Clifford and Croker, 1987), general location, name (optional), and feature class. The life-span can be an open or closed interval. An open interval represents all known time and applies to entities with unknown creation and termination dates. Mountains or lakes are examples of geographic entities with this type of life-span. Partially open intervals are those with either an unknown creation date or an unknown termination date. A building can be an example of a geographic entity with such a life-span as its creation date is known and it is expected to be in existence for an unknown duration. A closed interval indicates an entity with a known existence but that is no longer in existence. A building can be an example of a geographic entity which may have a closed interval life-span indicating that both its construction and destruction dates are known. The life-span provides the temporal bound for queries on a geographic entity against the information sources. In other words, a search would not be performed for geographic entity representations on information sources pre-dating the entity's life-span start date or post-dating its life-span end date.

Each geographic entity has a general spatial location referred to as the geographic entity footprint. A spatial location representation reflects varying degrees to which it is known. If the geographic entity is highly mobile, it can have an open-ended spatial location comparable to the open-interval life-span in that its bounds are unknown. The open interval is assumed to be the extent of the Earth's surface. If the bounds of an entity's movements are known (for example, an animal's habitat range), the footprint is a region defining the bounds of known (or predicted) movements. Fixed geographic entities will have as a footprint the minimum bounding rectangle of all recorded positions. Comparable to the role of the life-span, the footprint sets the spatial bounds for queries for geographic entities against the information store. Figure 4 summarizes the minimum set of attributes for a geographic entity.

The Geographic Entity Register is linked to registers describing specific independent components of spatiotemporal change. These registers link instances of geographic entities with the information sources in which they are depicted. In our case, as mentioned above, we assume three independent components of change for an object, resulting in three registers:

- Boundary register, containing information on outlines (e.g., edges of a geometric entity);
- Position register, describing the position of the object (e.g., *X*, *Y*, *Z* coordinates and orientation information); and
- Thematic register, describing non-spatial characteristics of a geographic entity (e.g., identifying building as residential).

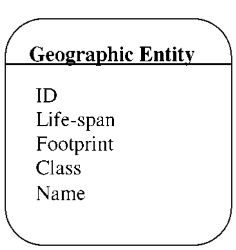

Figure 4. Attributes for a geographic entity.

Boundary Register

The Boundary Register contains boundary representation objects that describe a relation between a geographic entity and an individual information object from the information store. The collection of boundary representation objects for one geographic entity tracks the multiple spatial configurations of that geographic entity. Each boundary representation of a geographic entity is extracted from a specific information source. In this case, only spatial information objects contribute meaningful boundary representations. For example, a map in the information store can produce a detailed two-dimensional (2D) representation of a lake, a satellite image may contain a coarse 2D representation of the same lake, and a stereomodel can produce a detailed 3D representation of a building. The boundary register maintains (Figure 5)

- a unique boundary ID,
- a geographic entity ID,
- an information object ID (linking to the source of the information),
- the time stamp of the information object,
- information on the extraction method (specifically accuracy measures for the extracted object), and
- the spatial representation of the geographic object as extracted from the information object (e.g., vectorized or raster outlines, with or without accompanying radiometric information).

The boundary description of a geographic entity as extracted from an information source can take a number of forms. Typically, the description is an ordered set of points, lines, and regions depending on the complexity of the entity and the detail of the information source. In the case of imagery, a subset of the image is extracted and stored along with the extracted entity outline. This image subset is valuable because it contains radiometric data that can be used for quick subsequent reprocessing of the entity outline. Access to the radiometric data is useful in cases where it becomes necessary to revisit prior information, e.g., when contradictory information raises suspicion of errors in a previous boundary extraction process.

Position Register

The Position Register contains a set of position representation objects that describe a relation between a geographic entity and an information object. A position representation object has many of the same attributes as a boundary representation object. It includes a unique ID, the geographic entity ID, an information object ID, the information object time stamp, the position extraction method, and the position description (see Figure 6). The position description is most typically an (*X*, *Y*, *Z*) set of coordinates of the center of the geographic entity's minimum bounding rectangle (MBR).

B.Position: $(X, Y, Z, \theta)_{IO_{ti}}$ describes B's location in the information object IO_{ti} with time stamp ti; *X*, *Y*, and *Z* are the

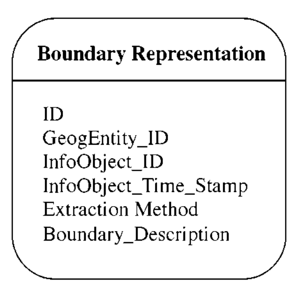

Figure 5. Attributes of a boundary representation object.

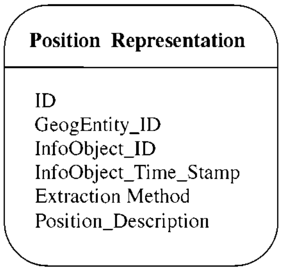

Figure 6. Attributes for a position representation object.

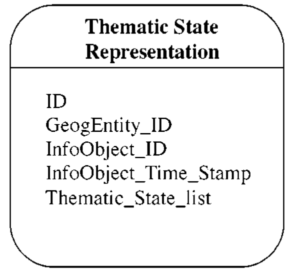

Figure 7. The minimum set of attributes for a thematic state representation object.

coordinates describing the MBR center, and θ is an initial azimuth for the object if a record of azimuth is pertinent.

The position descriptions are the key primitives for detecting movement. Movement is defined as a translation or a rotation about an axis (a definition similar to that used by Moreira *et al.* (1999)). Movement of an object can not be detected from a single information source (a single frame in the case of video) because a single source is assumed to represent an instant of time. The detection of movement requires at least two information sources with different time stamps defining a time interval. A method on two consecutive time ordered position representation objects for the same geographic entity yields a unit movement. We define a unit movement as the set of deltas between two position representations: $\text{B.movement}_{IO_1,IO_2} = (\Delta X, \Delta Y, \Delta Z, \Delta\theta)$

A record of movement consists of the unique ID of a geographic object, a distance and direction vector indicating movement of the geographic entity, the IDs of the two information objects from which the vector was extracted, and the time interval computed as the difference between the two information object time stamps.

Thematic State Register

The Thematic State Register maintains a collection of thematic state representation objects. These objects represent a relation between geographic objects and individual information objects. States are multidimensional, so a geographic entity may have several concurrent thematic states. A house, for example, may have the concurrent states: exists, has new owner, has new roof, and is in violation of a zoning ordinance. An individual information object can contribute information on one or more states. Information in the Thematic State Register consists of a unique ID, the ID of the geographic entity being described, the ID of the information object from which the states were extracted, the time stamp of the information object, and the states of the geographic entity contributed by the information source (see Figure 7). Some information sources (sensors) are designed to record a change of state the moment it occurs (e.g., a flood gauge), in which case the time stamp captures the instant of change. Other information sources can document changes of state but not at the moment they occur, which is typically the case for satellite imagery. An image can record that a house burned down but not necessarily when. The IO__time stamp only indicates that the state was in effect at the time stamp date.

Image-Based Change Detection within the Gazetteer

Description of the Approach

The identification of changes in geographic entity outlines is a fundamental operation within our model. The method we are presenting here employs least-squares matching for the detection and tracking of edges in digital images. Our approach is along the lines of Gruen (1985), Gruen and Stallman (1993), and Gruen and Agouris (1994), with some modifications introduced to accommodate desired resolutions of change monitoring, and to better handle noise through the analysis of local edge information within the matching windows.

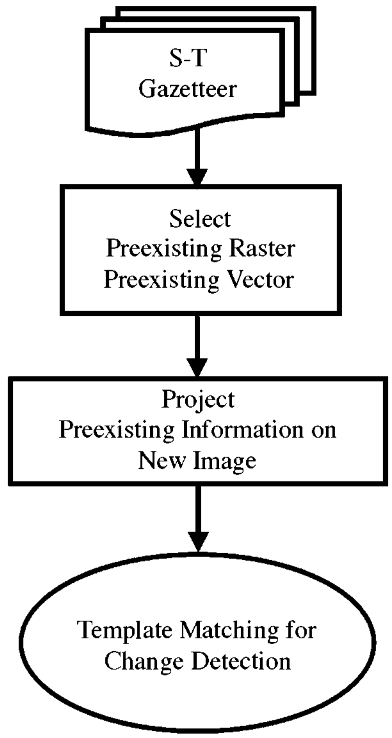

Figure 8. Process flowchart for outline change detection.

Our change-detection process is outlined in Figure 8. Essentially, we proceed by using a window depicting an edge pattern as a reference template, and matching this window to edge locations in a new image. By selecting an edge template depicting pre-existing information on a geospatial object (extracted from the content of Gazetteer registers), and matching it to a new image depicting this object at a newer instance, we are able to identify whether this object has remained unchanged or not.

Pre-existing boundary and position information is retrieved from our SpatioTemporal Gazetteer. Using this information for the vector and georeferencing information for the new image, we transfer the old object boundary onto the new image. There, it is used as input for a differential application of least-squares template matching (DLSM). Older information in raster format produces object outline (edge) templates that are matched to the image. By matching an edge template to the new image, we identify outline locations in this new image as conjugates to template edge locations. Thus, we transfer the high accuracy potential of least-squares matching onto object extraction. By performing this process at select positions along the object outline, we compare the complete object to its last recorded boundary. Changes are identified as instances where DLSM fails to match adequately.

Object Decomposition

Our algorithm for change detection follows a point matching approach, with prior information applied on a new image. A problem arises during this process related to the points that would sufficiently represent the 3D object on the 2D image space. To compensate for that, a 3D object is expressed through a wireframe representation (de Cambray, 1993; Vosselman and Veldhuis, 1999), based on prior vector information.

Considering change detection in building outlines, the examined object may be considered as an aggregation of planar surfaces, following the concept of polyhedral models (Bignone *et al.*, 1996; Hendrickx *et al.*, 1997). A generalization of planar surfaces may be performed by assuming that they are equivalently represented by intersections of planes, i.e., by lines. Due to the nature of our raster dataset (aerial photography), vertical aerial photos are assumed to be available, and 3D planes are merged into two dimensions, by applying an overlay operation on our 3D vector based on our viewpoint perspective. While building outlines tend to be geometrically regular shapes, irregular outlines (e.g., outlines of landlots and geographic regions) can be approximated by polygonic lines, and are behaving similar to regular outlines over brief intervals.

To assure an adequate representation of lines from points, a new element is introduced, the Minimum Spatial Element (MSE) (Mountrakis *et al.*, 2000). The MSE describes the resolution of spatial change in which the user is interested. Absolute estimates for this resolution (e.g., 2 m, or 0.5 pixel), or relative measures (e.g., the average size of a room) can be assigned. By using this information, we perform a segmentation of outlines, and lines are essentially substituted by the corresponding points along the outline. As corners are defined by line intersections, we do not have to consider them in our outline decomposition. A final product of this process is moving from a 3D object to a reduced set of points in the 2D image space (Figure 9).

Mathematical Model

Our change-detection method employs least-squares matching for the detection and tracking of edges in digital images. A window depicting an edge pattern is introduced as a reference template that is subsequently matched to digital image patches in the vicinity of actual edge segments. The concept behind the method is simple yet effective: by matching the edge template window to an image window, we can identify edge locations in the image as conjugate to the *a priori* known template edge positions (Gruen and Agouris, 1994).

Assuming $f(x,y)$ to be the reference edge template and $g(x,y)$ to be the actual image patch, a matching correspondence is established between them when

$$f(x, y) = g(x, y). \tag{1}$$

However, considering the effects of noise in the actual image, the above equation becomes

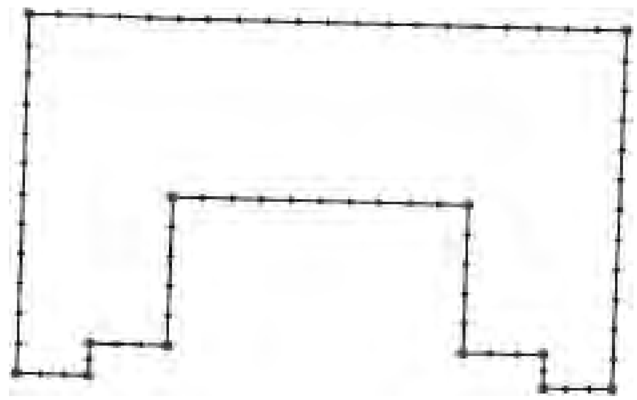

Figure 9. Identifying check points along the outline of a building.

$$f(x, y) - g(x, y) = e(x, y) \tag{2}$$

with $e(x,y)$ being the error vector.

In a typical least-squares matching method, observation equations can be formed relating the gray values of corresponding pixels, and they are linearized as

$$f(x, y) - e(x, y) = g^0(x, y) + \frac{\partial g^0(x, y)}{\partial x}dx + \frac{\partial g^0(x, y)}{\partial y}dy. \tag{3}$$

The derivatives of the image function in this equation express the rate of change of gray values along the x and y directions, evaluated at the pixels of the patch. Depending on the type of edge, the geometric relationship describing the two windows may be as complex as an affine transformation, or as simple as a simple shift and/or rotation. Regardless of the choice of geometric transformation, the resulting observation equations are grouped in matrix form as

$$-\mathbf{e} = \mathbf{Ax} - \mathbf{l}; \mathbf{P}. \tag{4}$$

In this system, $\mathbf{l}$ is the observation vector, containing gray value differences of conjugate pixels. The vector of unknowns $\mathbf{x}$ comprises the shift at the x direction, while $\mathbf{A}$ is the corresponding design matrix containing the derivatives of the observation equations with respect to the parameters, and $\mathbf{P}$ is the weight matrix. A least-squares solution allows the determination of the unknown parameters as

$$\hat{\mathbf{x}} = (\mathbf{A}^T\mathbf{PA})^{-1}\mathbf{A}^T\mathbf{Pl}. \tag{5}$$

While the above formulas reflect a standard template matching method, our problem introduces certain challenges. Indeed, comparing templates of the same object in various time instances and often captured by different cameras introduces large amounts of noise. In order to optimize the performance of our template matching method, we have to minimize the effect of radiometric variations among the two images (e.g., due to noise, differences in general histogram properties, or even different resolutions).

Weight Manipulation to Minimize the Effects of Noise

The differences in general exposure conditions among two different images in two distinct time instances may affect substantially the performance of the above described matching method. Considering, for example, the ledge of a roof, its appearance may be affected by variations in the two missions

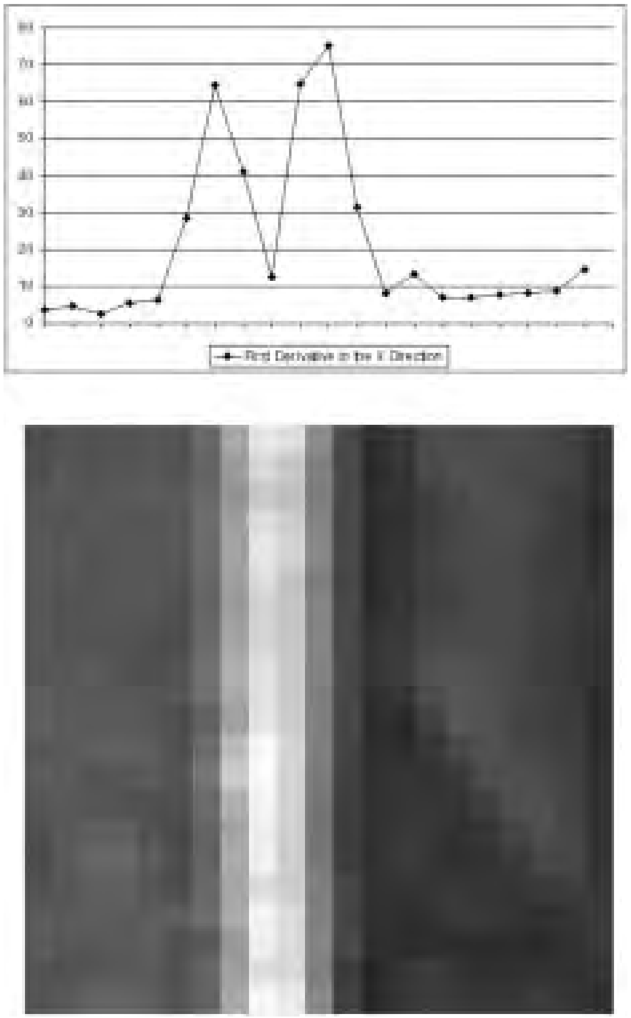

Figure 10. Template window (top) and corresponding averages of the first derivative in the *x* direction (bottom).

that produced the processed images. However, while individual gray values may differ, the main information content of the scene, namely, the outline of the ledge, remains the substantial piece of information that allows us to establish correspondences between two conjugate patches.

In order to minimize the effects of variations on our solution, we have to allow edge pixels to influence the solution of Equation 5 above more than the rest of the template. This can be performed by manipulating the corresponding weight matrix **P** (of Equations 4 and 5). Indeed, by assigning high weight values to certain pixels, we allow them to influence the solution more than pixels with weight values approaching 0. Accordingly, we enhance the solution of the model presented earlier by incorporating local edge analysis in it. This transforms our template matching from a common area-based matching scheme to a content-based matching process, improving its performance potential.

Considering, for example, the window of Figure 10, we calculate the average first derivative of gray values along the *x* direction: i.e.,

$$g_x(x, y) = \frac{\partial g(x, y)}{\partial x}\,; \tag{6}$$

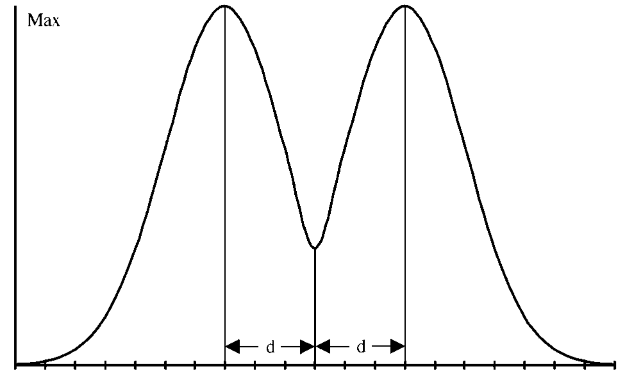

Figure 11. Weight assignment along the *x*-axis for the edge set-up of Figure 10.

$$\bar{g}_x(x) = \frac{\sum_{1}^{y} g_x(x, y)}{y}\,. \tag{7}$$

An analysis of the $\bar{g}_x(x)$ graph (Figure 10), by using the first and the second derivatives, reveals three maxima, two of which are substantially stronger than the remaining one. The dominant maxima correspond to the outline of the depicted roof ledge.

Two criteria are introduced in the decision process to accept or reject a maximum that in essence corresponds to an edge. If we define $\max_0$ as the maximum of the three maxima found: i.e.,

$$\max_0 = \mathrm{Max}\{\max_1, \max_2, \max_3\},$$

then we compare this value with the existing maxima ($\max_I$): i.e.,

- If $\max_i >$ HIGH% $\max_0$, then that edge is accepted because the high value of relative sharpness reveals a strong edge.
- If $\max_i <$ LOW% $\max_0$, then that edge is rejected because the low value of relative sharpness reveals a weak edge.
- If LOW% $\max_0 \le \max_i \le$ HIGH% $\max_0$, then we compare the second derivative, before and after the candidate maximum, by using the criterion

$$|\bar{\nabla}_x(x)| + |\bar{\nabla}_x(x+1)| \le \text{HIGH\% } \max_0.$$

If the above condition is satisfied, then a point is rejected; otherwise, it is included in the analysis. The above-mentioned HIGH and LOW percentage values may be viewed as general parameters of our approach. In our experimental results we have employed values ranging up to 30 percent for the LOW% and down to 70 percent for the HIGH%.

With the two above criteria, a "relative" check based on maximum value compares the dominant edge with other gradient maxima to eliminate false responses. This is performed along various directions (selecting orientation steps, e.g., 30 degrees, and repeating the process). In essence we are pursuing linear edges within our windows. We use Gaussian distributions to describe the reduction of weight coefficients as we move away from an edge.

Once we identify these linear edge segments, we assign higher weights to their corresponding pixel locations (Figure 11). This allows us to incorporate window geometry information in the matching solution.

Implementation Considerations and Experimental Results

In order to test the change-detection algorithm performance, a prototype SpatioTemporal Gazetteer was created for a semi-urban scene (campus of the University of Maine). Our multimedia sources provide a description of the campus during the last century. The available datasets include

- aerial imagery datasets representing the campus over the last half of the century (period 1949–1997),
- corresponding vector datasets of extracted outlines on the above images, and
- a variety of thematic datasets (e.g., building blueprints, building usage records).

The prototype implementation was performed using Microsoft's Visual Studio on a stand-alone PC.

To illustrate information flow within the Gazetteer for change-detection purposes, let's consider a new image introduced in the Multimedia Information Store. As it becomes an information object, its footprint (e.g., derived by its orientation parameters) is identified and stored as attribute information. Using the image footprint as a spatial target, we search the Geographic Entity Register for any entities whose footprints overlap the image footprint. This matching operation uses the eight region-region spatial relations defined by Egenhofer and Franzosa (1991). A list of IDs is returned for any geographic entities whose footprints are contained by, covered by, or overlap the image footprint. Using the returned list of entity IDs, we search the boundary and position (and even thematic if that is of interest) registers for representations of these entities. Boundary and position representations returned from the corresponding registers are input for the boundary change-detection tool. The last recorded boundary representation is projected onto the new image and compared using the matching technique. If the object has not changed since its last record in the STG, the time stamp of all representations is updated in the corresponding registers. If our differential matching algorithm detects changes in the object, the corresponding attributes (namely, extraction method and boundary description in the boundary register, and extraction method and position description in the position register) have to be updated. A comparison of the detected change patterns to common sensor error models is performed to distinguish sensor errors (e.g., misregistration) from actual change (Mountrakis, 2000). Accepted change information is used to update the corresponding Gazetteer registers. At the same time, the time stamp information and info-object ID in both registers are updated to reflect the time indicated by the processed image and the ID of this image. Figure 12 is a representation of this process.

From a practical point of view, the separation of boundary change and movement presents an interesting challenge. The minimum-bounding rectangle (MBR) and the center of gravity are two of the most common methods used to indicate the position of two-dimensional objects. As an object changes its boundary, both of these indices change, implying movement, whereas we might have a simple expansion of it. A solution to this problem is offered by examining the number of points that remain unchanged during our matching process. If an adequate percentage of sequential points remain fixed, we recognize that the object has not moved, and that only a part of it has been modified. In our approach this percentage is selected empirically, and it typically ranges between 25 percent and 50 percent of the points (selected to represent one to two sides of a four-sided object).

In order to examine the performance of our change-detection tools, we manually extracted buildings from older aerial images, and overlaid them on subsequent raster images using relevant orientation information. Using this approach, we performed change detection in order to capture the evolution of buildings during this last century. Several tests were performed to ensure the reliability of our algorithm. Our main objective was to correctly identify changed areas, and to verify unchanged points. Three factors were considered through the evaluation:

- Good Detection: the ability to locate and mark all the real edges;
- Good Localization: minimal distance between the detected edge and real edge; and
- Clear Response: only one response per edge.

Good localization is the major advantage of LSM. Indeed, experiments showed that the extracted edges approximated the real ones at subpixel accuracy (approx. 0.1 pixel). For the other two factors, a significant improvement was noticed when the edge levels were three or more and we were able to establish a geometric representation of the edge. This geometric analysis proved to be very fast because the weight matrix had to be constructed once for every matching point. The width of the level(s) establishes a scaling factor in the whole process, which guides the template quickly and accurately to the new edge, when there is no change. Within our testing, almost 100 percent accuracy was achieved when edge geometry was incorporated in the solution through weight manipulation. Otherwise, random noise such as building windows, shadows, or cars affects the algorithm. Edge geometry analysis allows us to distinguish random noise from the real edge through the manipulation of the weight matrix. Especially in the case of shadows, a problem inherently difficult because the artificial edges are very strong, the expected geometry of the edge, expressed through the size of the curb, introduced the essential metric reasoning to compensate for such errors. Figure 13 shows typical results from the application of our matching tool within the SpatioTemporal Gazetteer established for the University of Maine campus. The results indicate that the boundary of the building has changed. However, because only few points have changed, we recognize that this change is only on the boundary, and does not imply a movement of the object.

Conclusions and Future Work

The SpatioTemporal Gazetteer framework we presented addresses the need for rapid access to up-to-date precise geospatial object information for modern decision making. It makes use of multiple information resources and incorporates components to track changes to objects over time. This Gazetteer enables the management of representations of instances of geographic entities and their changes over time in an explicit manner, instead of the implicit handling of such information offered by common layer-based approaches. It allows us to organize geospatial information in an object-oriented manner that captures essential components of the spatiotemporal behavior of objects.

In order to ensure information flow within our gazetteer environment, we have modified template matching to perform in a multitemporal, multisensor environment. The presented algorithm extends the concept of least-squares template matching to automatically perform change detection in a spatiotemporal environment. By taking advantage of gazetteer information and metadata information, we can reduce the problem of 3D object monitoring to an image-space 2D matching problem. Furthermore, by manipulating the weight matrix, we can incorporate semantic information in our matching process, minimizing the effects of random noise and radiometric variations often associated with the analysis of images captured by different sensors at different time instances. While our experiments focused on vertical photography, the mathematical foundation of our algorithms would accommodate equally well oblique imagery, even close-range photos.

We plan to extend our semantic reasoning to incorporate scene analysis, as well as an analysis of the effect of sensor orien-

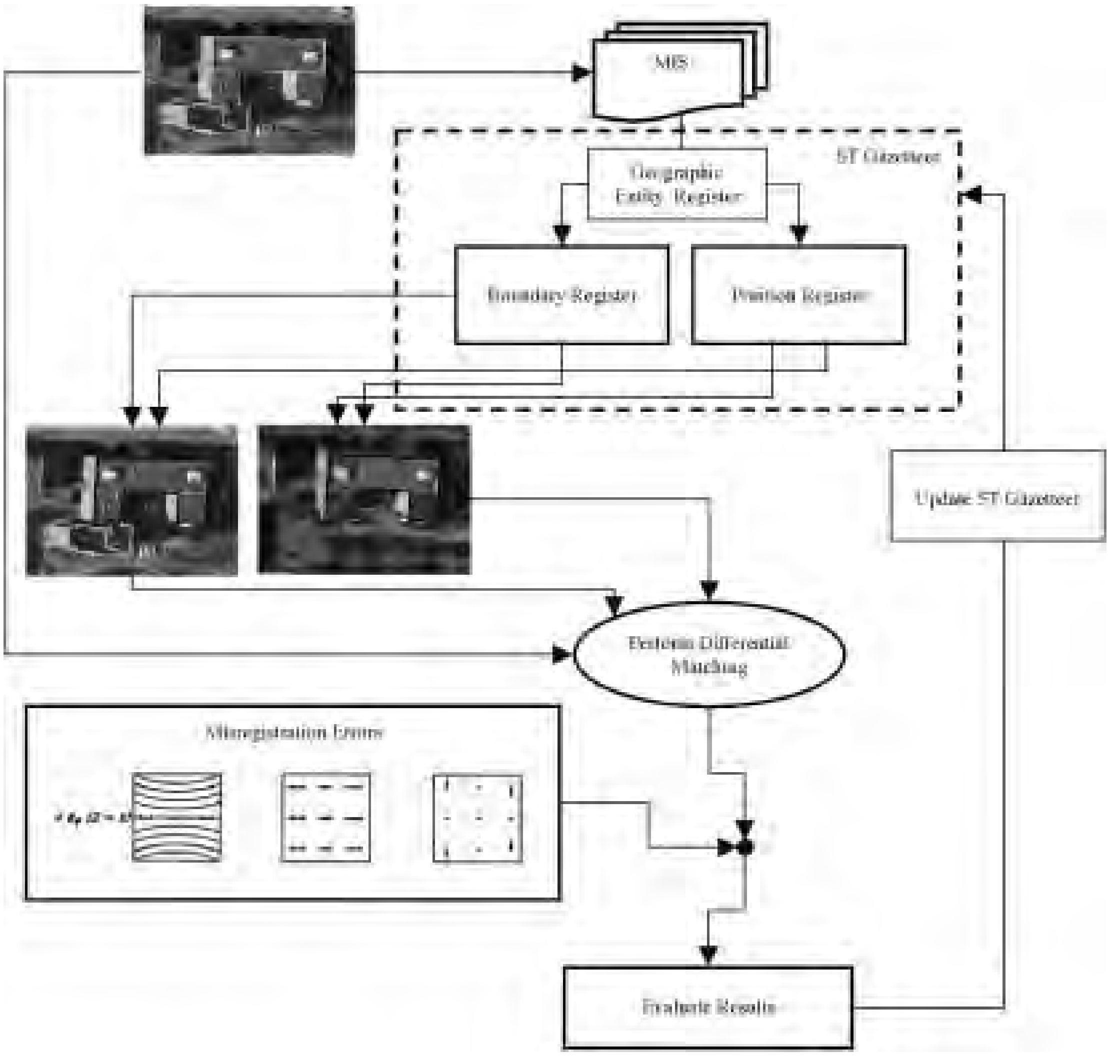

Figure 12. Information flow within the SpatioTemporal Gazetteer.

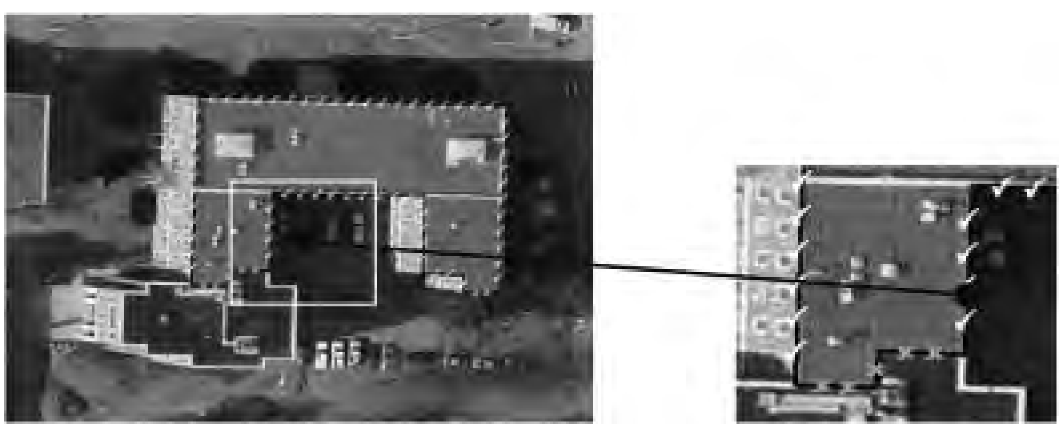

Figure 13. Example of our application environment, where *x* indicates no match (and therefore change), and √ successful match (and therefore no change).

tation errors in our change-detection process. Currently, our matching process is performed locally for an object, and independently of its broader scene. We plan to incorporate scene analysis metrics to tie the change-detection processes of individual buildings into a larger, scene monitoring process. This will allow us to perform complex scene analysis tasks, and to detect and eliminate the effect of systematic errors (like erroneous orientation data and misregistration) on our change-detection process. We will also investigate interoperability issues related to the deployment of a network of distributed Gazetteers.

Acknowledgments

This work was supported by the National Imagery and Mapping Agency through NURI grant number NMA202-98-1-1113, and by the National Science Foundation through CAREER grant number IIS-9702233. We would also like to thank Mr. Mustafa Palancioglou for his valuable comments and contributions to this research project.

References

Batty, M., and P. Densham, 1996. *Decision Support, GIS, and Urban Planning*, Centre for Advanced Spatial Analysis, University College London, http://www.geog.ucl.ac.uk/~pdensham/s_t_paper.html.

Beard, M.K, and T.R. Smith, 1998. A Framework for Meta-Information in Digital Libraries, *Managing Multimedia Data: Using Metadata to Integrate and Apply Digital Data* (A. Sheth and W. Klaus, editors), McGraw Hill, Inc., New York, N.Y., pp. 341–365.

Beard, M.K., T.R. Smith, and L. Hill, 1997. Meta-Information Models for Georeferenced Digital Library Collections, *Proceedings of Second IEEE Conference on Metadata*, 16–17 September, Silver Springs, Maryland, URL: http://computer.org/proceedings/meta97/papers/kbeard/kbeard.html.

Bignone F., O. Henricsson, P. Fua, and M. Sticker, 1996. Automatic Extraction of Generic House Roofs from High Resolution Aerial Images, *Computer Vision - Proceedings of ECCV '96*, Lecture Notes in Computer Science Vol. 1064, Springer-Verlag, Berlin, pp. 85–96.

Clifford, J, and A. Crocker, 1987. The Historical Ralational Data Model (HRDM) and Algebra Based on Lifespans, *Proceedings IEEE Conference on Data Engineering*, 03–05 February, Los Angeles, California, pp. 528–537.

de Cambray, B., 1993. Three-Dimensional Modeling in a Geographical Database, *Auto Carto 11*, 30 October–01 November, Minneapolis, Minnesota, pp. 338–347.

Del Bimbo, A., E. Vicario, and D. Zingoni, 1995. Symbolic Description and Visual Querying of Image Sequences Using Spatio-temporal Logic, *IEEE Transactions on Knowledge and Data Engineering*, 7(4):609–621.

Egenhofer, M., and R. Franzosa, 1991. Point Set Topological Relations. *International Journal of Geographic Information Systems*, 5(2):161–174.

Goodchild, M., 1998. The Geolibrary, *Innovations in GIS 5* (S. Carver, editor), Taylor and Francis, London, pp. 59–68.

Gruen, A., 1985. Adaptive Least Squares Correlation: A Powerful Image Matching Technique, *South African Journal of Photogrammetry, Remote Sensing & Cartography*, 14(3):175–187.

Gruen, A., and P. Agouris, 1994. Linear Feature Extraction by Least Squares Template Matching Constrained by Internal Shape Forces, *International Archives of Photogrammetry & Remote Sensing*, 30(Part 3/1):316–323.

Gruen, A., and D. Stallmann, 1993. High Accuracy Edge Matching with an Extension of the MPGC-Matching Algorithm, *Applications of Geodesy to Engineering*, AG Symposia Vol. 108, Springer-Verlag, Heidelberg, Germany, pp. 339–350.

Hendrickx, M., J. Vandekerckhove, D. Frere, T. Moons, and L. van Gool, 1997. On the 3D Reconstruction of House Roofs from Aerial Images of Urban Areas, *International Archives of Photogrammetry and Remote Sensing*, 32(Part 3–4w2):87–96.

Hill, L., J. Frew, and Q. Zheng, 1999. Geographic Names: The Implementation of a Gazetteer in a Georeferenced Digital Library, *D-Lib Magazine*, Vol. 5, No. 1, http://www.dlib.org/dlib/january99/hill/01hill.html.

Hornsby, K., 1999. *Identity Based Reasoning about Spatio-Temporal Change*, PhD dissertation, University of Maine, Orono, Maine, 164 p.

Lombardo, D., and Z. Kemp, 1997. Toward a Model for Mulitmedia Geographical Inforrmation Systems, *Innovations in GIS 4*, Taylor and Francis, London, pp. 56–69.

Moreira, J., R. Cristina, and J.-M. Saglio, 1999. Representation and Manipulation of Moving Points: An Extended Data Model for Location Estimation, *Cartography and Geographic Information Systems*, 26(2):109–123.

Mountrakis, G., 2000. *Image-Based Change Detection Using an Integrated Spatiotemporal Gazetteer*, M.Sc. Thesis, University of Maine, Orono, Maine, 149 p.

Mountrakis, G., P. Agouris, and A. Stefanidis, 2000. Navigating through Hierarchical Change Propagation in SpatioTemporal Queries, *Proceedings TIME 2000*, 07–09 July, Cape Breton, Nova Scotia, Canada, pp. 123–131.

NIMA (National Imagery and Mapping Agency), 2000. *The Big Idea Framework*, National Imagery and Mapping Agency, Bethesda, Maryland, URL: http://www.opengis.org/thebigidea/.

Salzberg, B., and V.J. Tsotras, 1999. A Comparison of Access Methods for Time-Evolving Data, *ACM Computing Surveys*, 31(2):158–221.

Vosselman, G., and H. Veldhuis, 1999. Mapping by Dragging and Fitting of Wire-Frame Models, *Photogrammetric Engineering & Remote Sensing*, 65(7):769–776.