Hurricane Disaster Assessments With Image-Driven Data Mining in High-Resolution Satellite Imagery

Christopher F. Barnes, *Senior Member, IEEE*, Hermann Fritz, and Jeseon Yoo

Abstract—**Detection, classification, and attribution of high-resolution satellite image features in nearshore areas in the aftermath of Hurricane Katrina in Gulfport, MS, are investigated for damage assessments and emergency response planning. A system-level approach based on image-driven data mining with σ-tree structures is demonstrated and evaluated. Results show a capability to detect hurricane debris fields and storm-impacted nearshore features (such as wind-damaged buildings, sand deposits, standing water, etc.) and an ability to detect and classify nonimpacted features (such as buildings, vegetation, roadways, railways, etc.). The σ-tree-based image information mining capability is demonstrated to be useful in disaster response planning by detecting blocked access routes and autonomously discovering candidate rescue/recovery staging areas.**

Index Terms—**Emergency response planning, image-driven data mining, image information mining, satellite image hurricane disaster assessments, σ-tree classifiers.**

I. INTRODUCTION

HIGH-RESOLUTION multispectral commercial satellite imagery is available for most locations on Earth before and after major disasters. Potential uses of these image sets include the following:

- debris field detection;
- ingress route obstruction detection;
- building damage detection and assessments;
- remote sensing in tactical support of rescue planning.

Each of these image exploitation applications for remotely sensed images can be performed manually with expert human analysis and image annotation tools, but such effort is labor intensive and hinders the quick response needed for first responders in large disaster impact zones. Tools that are able to adapt to newly encountered disaster-related features and that can be easily trained and applied to process volumes of imagery of large disaster regions are needed. This paper explores such a tool that is based on image-driven data mining with σ-trees [2]. This data mining approach to feature extraction and image content labeling is adaptable to various types of sensor systems and data sets, and enables extraction algorithms to be configured on-the-fly by novice users. In these facets, the demonstrated framework is adaptable to nearly any disaster-related image exploitation task without requiring time and expense in pattern recognition engineering efforts.

The data mining image assessment tools demonstrated in this paper are able to support interactive learning sessions performed under the tutelage of *sensei users*, who are granted rights of data annotation for training purposes in the data warehouse. These data mining tools are designed to be easy to use (do not require a high level of end-user expertise in the underlying technologies) and are not time consuming or resource demanding to train (training can be performed with mouse-driven actions on laptop class computers). The supporting σ-tree feature exemplar templates are usually well generalized even when trained with limited training data. Hence, in this system, a sensei user can conduct an interactive learning session on a small study area, and then the computer-assisted assessment tool applies the lessons learned from the localized study area to autonomously analyze larger-area disaster impact zones.

A. Integrative Data Warehouses and Hard Truth Data

Much knowledge base value is created in the archival of analyzed imagery. Knowledge bases should be useful in supporting referential comparisons with new imagery for inferential decision support [5]. These warehouses may also incorporate "before disaster" feature ground truth data and, if available, "after disaster" ground truth data acquired in postdisaster site visits. This class of validated ground truth data is referred to as *hard truth* data in this paper.

Hurricane Katrina struck coastlines along the northern Gulf of Mexico particularly vulnerable to storm surge impact and flooding. The widespread failure of tide gauges along the Mississippi and Louisiana shores called upon reconnaissance crews to collect high water marks. The second author led a Georgia Tech reconnaissance team that surveyed impacted coastlines in Florida, Alabama, Mississippi, and Louisiana in September and October 2005. Four separate boat expeditions that covered the Gulf islands from Petit-Bois Island to Cat Island of Mississippi were also made. The teams measured maximum storm tides, overland flow depths, and inundation distances. The teams also collected soil samples from storm deposits and documented erosion (site visit photographs shown in Fig. 1). Perishable infrastructure damage was recorded at various scales. The elevations of flood water marks on buildings, scars on trees, and rafted debris were measured as indicators of the maximum storm tide. High water marks were photographed and localized using Global Positioning System. Transects from the beach to high water marks were recorded with a laser range finder.

Manuscript received May 1, 2006; revised November 5, 2006. This work was supported in part by the National Geospatial-Intelligence Agency and in part by the National Science Foundation SGER program Award CMS-0553144.

C. F. Barnes is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Savannah, GA 31407 USA.

H. Fritz and J. Yoo are with the School of Civil and Environmental Engineering, Georgia Institute of Technology, Savannah, GA 31407 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TGRS.2007.890808

0196-2892/$25.00 © 2007 IEEE

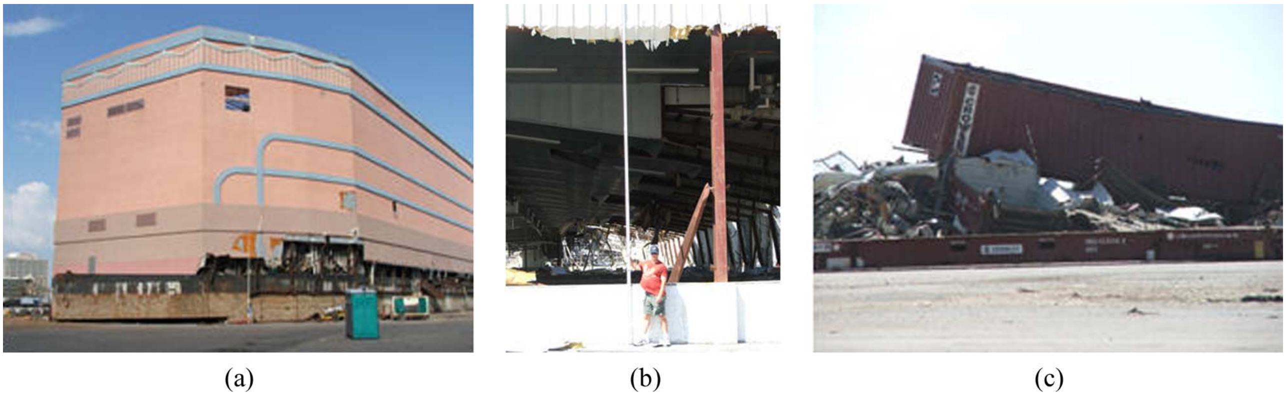

Fig. 1. Hurricane Katrina damage site measurements at Gulfport Mississippi terminal and container storage lot. (a) Casino barge washed ashore. (b) Surge wash-out height. (c) Sea and rail containers debris.

This survey showed that the storm tide peaked to the east of Katrina's path with consistent recordings between 7–10 m along a 60-km stretch of the Mississippi coastline from Lakeview (20 km east of center) to Ocean Springs. The state port of Mississippi in Gulfport is located 50 km to the east of the hurricane center in the middle of the hardest hit stretch of coastline. The surge penetrated at least 10 km inland in many portions of coastal Mississippi and up to 20 km inland along bays and rivers, crossing Interstate 10 in many locations. The high water marks dropped below 5 m along the Alabama coast. More than 2-m high water marks were measured 240 km east of Katrina's track along Florida's panhandle. The barrier islands 10–20 km offshore the Mississippi and Alabama coast were completely submerged and overwashed by the storm tide. The variation of the high water measurements on the islands was in accordance with corresponding onshore recordings.

B. Interactive Learning Sessions and Soft Truth Data

Disaster assessment applications need flexible computer tools that are not brittle with respect to a restricted set of predetermined disaster-related feature classes. Sensor modality issues such as time of day, weather state, and scene location specifics tend to make referential mining systems brittle with respect to static warehouses that totally depend on predisaster image collections for knowledge base formation. Assessment applications should employ adaptable tools capable of dealing with newly acquired image sets with feature (disaster related or not) classes not necessarily known or annotated beforehand in the warehouse. This capability is achieved in the proposed system of this paper with simple training steps performed under the supervision of a sensei user. The labeling activities of a sensei user in an interactive learning session create another class of "truth" data resident in a data warehouse. This class of subjectively labeled data is termed here as *soft truth* data.

C. Related Image Information Mining Research

Data mining aims to extract useful information from large archives of raw numerical data [6]. Image data mining concepts are still being formulated, and various systems have been proposed [10]. To some, "image mining" is the process of finding useful images in large image archives. To others, image mining implies the use of data mining in heterogeneous databases [7]. Query formulation domains are often restricted to only archived image metadata tags recorded in text or numeric formats [3]. Such metadata queries require annotation tags assigned beforehand to label image characteristics and feature content. Manual image annotations with a descriptive lexicon are prohibitive for high-volume workflows or large archives and are difficult when targeted toward disaster damage assessments because of lack of extensive relevant database image sets before specific disaster occurrences. Research has yielded image mining approaches such as content-based image retrieval, which is characterized by the ability of a system to retrieve relevant images based on an image's high-level query features such as color, texture, or morphology [9]. These features are all extracted by (sometimes compute-intensive) image analysis to enable content-based query rather than by using tagged metadata keywords [8]. One problem of high-level feature-to-feature comparative systems is the inability to adapt to new target feature sets and changing user needs, and being overly application specific in architecture and operation [5]. High-level feature-to-feature comparative systems have the challenging requirement for nonrecurring engineering for discriminating feature discovery and definition, and the development of algorithms for efficient feature extraction. The adaptability of such systems is constrained when encountering new sets of disaster imagery.

A novel template trellis called σ-tree is at the heart of the image information mining system utilized in this paper and allows queries to be performed in a data-compressed representation of the pixel space of the archive imagery. This system is easily adapted to new image sets and new features of interest. The next section describes, at a high level, σ-trees and their use for semi-autonomous and autonomous image content labeling. A more detailed description can be found in [2]. The demonstration and evaluation of σ-tree image-driven data mining on electro-optical satellite images to detect the highly specific features sets associated with hurricane disasters is the primary contribution of this paper. This paper quantifies the performance of image-driven data mining for image content labeling in hurricane damage assessments and for first-responder tactical rescue planning.

II. SOLDIER BIT-PLANE DATA MINING

A compressed binary data format useful in pattern recognition and data mining applications is the expanded digital "bit-plane" representations $\mathbf{i}_P$ formed by searches of application-specific σ-tree structures [2]. The σ-tree pattern recognition template set provides easily searched (structured) templates $\hat{\mathbf{t}}$ that approximate (unstructured) warehouse feature exemplars $\mathbf{t}$, and provides easily manipulated approximations of the content of new images being compared with the knowledge base. A normalized mean-squared distance measure $d(\mathbf{t}, \hat{\mathbf{t}}(\mathbf{i}_p))$ is used in a greedy sequential search encoder to quantify pattern recognition similarity and produce the sequence of expanded digital representations $\mathbf{i}_P = (i_1 i_2 \cdots i_P)$ (see the Appendix for additional details). A progression of data mining query searches can then be performed by comparing the resulting σ-tree bit-plane representations of the new image content with the expanded digital bit-plane representations of the archived knowledge warehouse content. The matched data indexes $\mathbf{i}_P$ (with full or partially matched bit-plane depth) can then be subsequently used for data mining extraction of estimated feature class membership and estimated feature attribute values. Hence, image content labeling with image-driven data mining starts with the formation of σ-trees for sets of referential feature exemplars of application interest contained in the warehouse knowledge base. For example, as shown in Fig. 2(a), assume that a *feature locator table* $\mathbf{L}_a$ has been created by a sensei user and stored in a conventional database system. Each row of the feature locator table describes an archived feature exemplar by specifying a source image $\mathbf{I}_{\text{source}}$, the pixel location (X, Y) of the feature exemplar, a descriptive feature label $\mathbf{F}$, and any hard or soft attribute labels or attribute values $\mathbf{A}$ associated with the feature example. The locator data tuples can thus be described as

$$\mathbf{L}_a = [k, \mathbf{I}_{\text{source}}, X, Y, \mathbf{F}, \mathbf{A}] \tag{1}$$

where k is the primary key of the database table $\mathbf{L}_a$, and a is used to differentiate feature locator tables.

A pixel block extraction tool combined with a "training" query Q_u into $\mathbf{L}_a$ is used to form a *snippet set* $\mathbf{S}_{Q_u}$ of pixel block feature exemplars. This user-generated query Q_u can be generally described as

$$\begin{aligned} Q_u = [&\text{select feature exemplar references from } \mathbf{L}_a, \text{and} \\ &\text{extract corresponding pixel blocks from } \mathbf{I}_{\text{source}} \\ &(\text{where user defined query condition is satisfied})]. \end{aligned}$$

The resulting training snippet set is indicated as

$$\mathbf{S}_{Q_u}(b_w, b_h, b_d) = \{\mathbf{t}(k_{Q_u})\} \tag{2}$$

where $\mathbf{t}(k_{Q_u})$ are $b_w \times b_h \times b_d$ pixel blocks identified by the foreign key k_{Q_u} (primary key of $\mathbf{L}_a$) returned by the query $Q_u(\mathbf{L}_a)$. The parameters (b_w, b_h, b_d) describe the width, height, and depth of the 2-D (spatial-only) or 3-D (spatial-spectral) pixel blocks. As shown in Fig. 2, the snippet set $\mathbf{S}_{Q_u}$ is used next as a training set to design a σ-tree T_σ with algorithms described in [1]. After a σ-tree $T_\sigma(\mathbf{S}_{Q_u})$ has been produced to represent the feature exemplars of $\mathbf{S}_{Q_u}$, this same σ-tree is used to generate compressed versions $\mathbf{i}_p(k_{Q_u})$ of $\mathbf{t}(k_{Q_u}) \in \mathbf{S}_{Q_u}$, where $\mathbf{i}_p(k_{Q_u})$ are a complete set of partial and full $(1 \le p \le P)$ expanded digital representations. This encoded and compressed representation of the feature snippet set is added to the data warehouse and entered into database tables indicated as

$$T_\sigma(\mathbf{S}_{Q_u}) = [k_{Q_u}, \mathbf{i}_1, \mathbf{i}_2, \ldots, \mathbf{i}_P] \tag{3}$$

where $\mathbf{i}_1 = (i_1)$, $\mathbf{i}_2 = (i_1, i_2)$, $\mathbf{i}_3 = (i_1, i_2, i_3)$, and so on. Tables such as these joined with other warehouse ground truth tables provide *Source Optimized, Labeled, DIgital Expanded Representations* or, more simply, "SOLDIER" representations of feature-relevant knowledge-base content.

The SOLDIER index tables $T_\sigma(\mathbf{S}_{Q_u})$ of (3) joined via k_{Q_u} with the feature and attribute information columns of the locator table $\mathbf{L}_a$ of (1) provide a powerful resource for image-driven data mining. During the data mining runtime, an extractor algorithm takes pixel blocks $\mathbf{t}_{(x,y)}$ from locations (x, y) of an image under analysis. Then, the SOLDIER σ-tree indexing engine T_σ is applied to $\mathbf{t}_{(x,y)}$ to yield the expanded digital representations $\mathbf{i}_p(x, y) = T_\sigma(\mathbf{t}_{(x,y)})$. The σ-tree-indexed representation $\mathbf{i}_p(x, y)$ of the pixel query block $\mathbf{t}_{(x,y)}$ can then be used in a conventional query to perform a data mine for feature similarity sets and aggregates of related label and attribute information, i.e.,

$$\begin{aligned} Q_{(x,y)} = [&\text{select } \mathbf{F} \text{ and } \mathbf{A} \text{ from } \mathbf{L}_a \text{ and } T_\sigma(\mathbf{S}_{Q_u}) \\ &\text{joined via } k_{Q_u}, \text{where } \mathbf{i}_p(x, y) = T_\sigma\left(\mathbf{t}_{(x,y)}\right) \\ &\text{equals } \mathbf{i}_p(k_{Q_u}) \text{ in } T_\sigma(\mathbf{S}_{Q_u})]. \end{aligned} \tag{4}$$

This SOLDIER index comparison between new image content and archived image content yields SOLDIER indexes with variable "bit-plane" index depth $p_{\max} \in \{0, 1, \ldots, P\}$. The reason for including partial SOLDIER representations is that partial σ-tree paths provide a sequence of pixel space generalizations (from course to fine) of the elements $\mathbf{t}(k_{Q_u})$ of $\mathbf{S}_{Q_u}$. This provides increased robust classifier performance when new image data are encountered that are only approximately similar, by varying degrees, to warehouse feature exemplars in the aggregate $\mathbf{S}_{Q_u}$. The query results of (4) will return sets $\{\mathbf{F}, \mathbf{A}\}$ of feature and attribute point data for each matched bit plane of the partial/full expanded digital representations $\mathbf{i}_p(x, y)$ of a pixel block $\mathbf{t}_{(x,y)}$. In addition to these mined set of feature and attribute data, it is possible to augment the data warehouse with archived maximum *a posteriori* (MAP) Bayesian estimates $\mathbf{F}_{\text{MAP}}$ and $\mathbf{A}_{\text{MAP}}$ over the data-mined feature/attribute sets $\{\mathbf{F}, \mathbf{A}\}$. The MAP values can be estimated from the training data in the following conventional manner. Let $N(\mathbf{i}_p)$ be the number of training vectors in the Voronoi cell of σ-tree exemplar $\hat{\mathbf{t}}(\mathbf{i}_p)$.

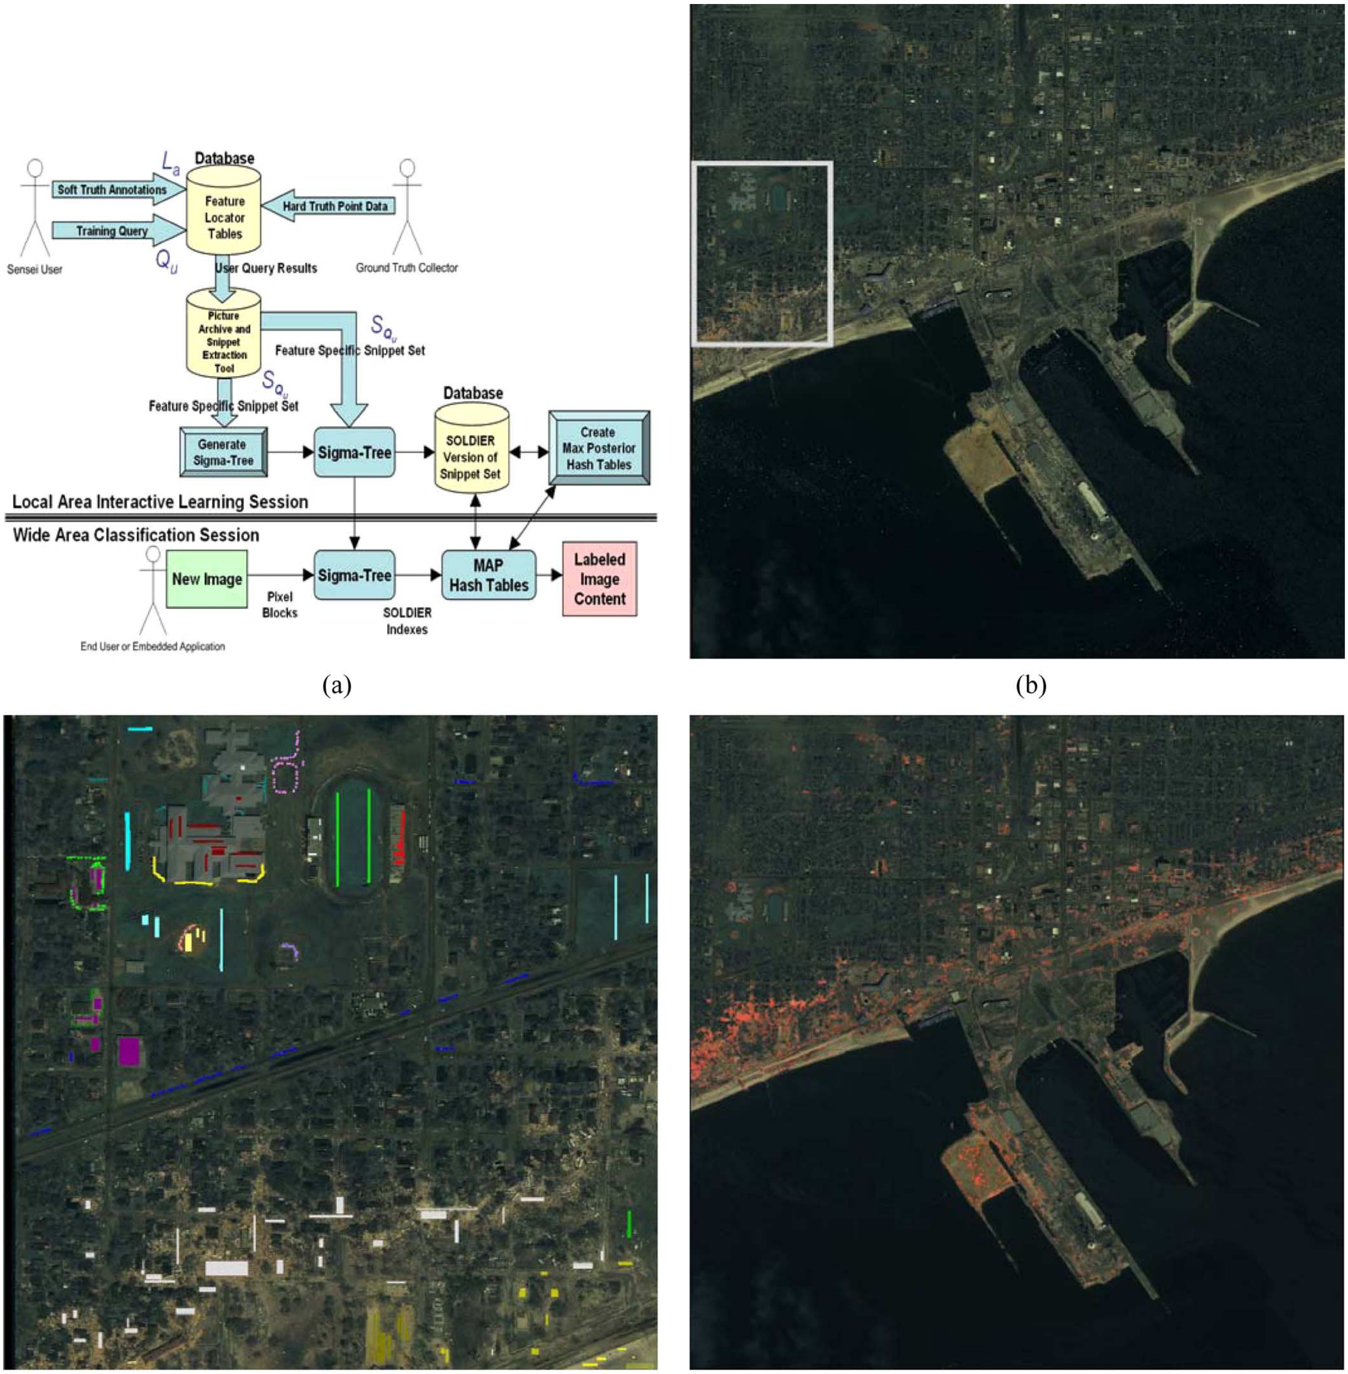

(a) (b)

(c) (d)

Fig. 2. Interactive training process and results for hurricane debris detection. (a) Image-driven data mining workflows for training and classification. (b) Local interactive study site selected by a sensei user for training a σ-tree hurricane debris field detection image data mine. (c) Sensei-user-selected feature (white) and counter-feature (colors) examples in the localized interactive training zone for hurricane debris detection σ-tree training. (d) Wide-area Hurricane Katrina debris detection result using the classifier from the sensei user's localized training session.

Let k_ω be the subset of the N exemplars that belong to class ω. Assuming equal prior probabilities over the class labels ω, the MAP estimate of $p(\omega|\mathbf{i}_p)$ is $k_\omega/N(\mathbf{i}_p)$. These MAP estimates are integrated into a sequence of hash tables $H_1[\mathbf{i}_1, \mathbf{F}_{\mathrm{MAP}}(1)],\ H_2[\mathbf{i}_2, \mathbf{F}_{\mathrm{MAP}}(2)], \ldots,\ H_P[\mathbf{i}_P, \mathbf{F}_{\mathrm{MAP}}(P)]$. An image-driven data mine search can be efficiently performed exterior to database services with the use of hash tables. That is, each bit-plane hash table $H_p[\mathbf{i}_p, \mathbf{F}_{\mathrm{MAP}}(p)]$ is easily utilized for efficient runtime MAP feature class labeling of mined image points by a search with the hash key $\mathbf{i}_p$ into $H_p[\mathbf{i}_p, \mathbf{F}_{\mathrm{MAP}}(p)]$ to return $\mathbf{F}_{\mathrm{MAP}}(p)$.

III. IMAGE-DRIVEN DATA MINING

The proposed image-driven mining system is able to support many image exploitation applications defined and configured by rather novice users in an easy, straightforward way. The following applications have been selected to demonstrate the simplicity and scope of potential applications of σ-tree-based image data mines when applied to remotely sensed disaster images.

A. Interactive Learning Session for Hurricane Debris Fields

The first disaster image mining experiment is the detection of nearshore hurricane debris fields in a situation where no prior archive imagery exists with such features, where no relevant warehouse-resident hard truth data exists, and where the only option for postdisaster satellite image exploitation is interactive learning on tutorial areas under the guidance of a sensei user. Shown in Fig. 2(b) is a commercial satellite image of the Gulfport, Mississippi area after Hurricane Katrina. This is a pan-sharpened multispectral composite formed from 2 × 2 m red, green, blue, and near-infrared planes (IKONOS multispectral) and 1 × 1 m panchromatic (IKONOS-2) imagery (creation date: September 2, 2005). The white box is the area selected by a sensei user for training in an interactive learning session. Shown near the bottom of Fig. 2(c) are small white boxes that represent image subareas with hurricane debris fields selected by the sensei user and labeled with a feature class label value $\mathbf{F}_{\text{Label}} =$ "Debris Field-Inundation Deposits." Each pixel location in each of the solid white boxes is archived and labeled. Then, the SOLDIER-MAP hash tables generated from this interactive training area are used for classification across the larger scene. A quick look test shows the sensei user in an interactive learning session that the following image content feature classes are sometimes confused with debris fields:

- Building-Roof-Gray;
- Building-Roof-Tan Gray;
- Building-Shadow;
- Stadium-Seats;
- Flood Swept-Tan;
- Flood Swept-Tan Gray;
- Flood Swept-Sandy;
- Parking-Asphalt;
- Parking-Concrete;
- Parking-Concrete-Edge;
- Road–Curb;
- Sidewalk;
- Grass;
- Soil-Red Clay.

The sensei user augments the warehouse-extracted training set during this interactive learning session by finding examples of each of these confuser feature classes, drawing a crop-box around such features (color boxes in Fig. 5), assigning the appropriate feature label, commanding the system to enter these feature exemplar pixel locations into a locator table, and repeating the experiment. The final result of performing an image-driven data mine for the entire scene against the sensei user's localized training session result is shown in Fig. 2(d). This result was obtained without any additional false alarm mitigation. An error analysis was performed to assess the classification error and specify the classifier's sensitivity as the *producer's accuracy* and the classifier's specificity as the *user's accuracy* [4]. The probability of an evaluation pixel being correctly classified is the producer's accuracy, determined by the number of pixels correctly classified divided by the total number of evaluation pixels for the class. Any shortfall indicates a measure of omission error probability. The probability that a pixel classified in the image is really a member of the assigned class is the user's accuracy and is calculated as the number of pixels correctly classified as a class divided by the total number of pixels that were assigned in that land cover class. The accuracy of this hurricane debris field feature extraction was assessed with sequestered test data not used for training. The evaluated user's accuracy for detecting hurricane debris is around 75%. The producer's accuracy is about 85%. These accuracy estimates were evaluated with about 1 480 000 test snippets. A comparative study was also conducted for debris field detection with a commercially available parametric multispectral classifier that uses maximum-likelihood decision rules. The parametric multispectral analysis yielded a sensitivity of 47% and a specificity of 82% when the following set of supervised classes were used for training: debris, building, house, road, grass, tree, soil, parking, sidewalk, and soil.

B. Land Cover and Disaster Impact Classification

The next experiment's objective is to classify and label relevant land cover content in the wide-area scene even if the relevant features are not significantly damaged or impacted by the hurricane disaster. The intent is to have an autonomous method of detecting where hurricane impact zones interrupt normal features that could support rescue/recovery logistics, e.g., debris fields that interrupt routes of ingress and egress. The prior list of features is expanded by a sensei user by adding exemplars to the warehouse of the following image content feature classes:

- Building-Roof-Edge;
- Building-Roof-Corner;
- Building-Roof-Damage;
- Building-Roof-Brown;
- Building-Roof-Dark Gray;
- Building-Roof-Damage;
- Ship Cargo Containers;
- Vehicle;
- Road–Asphalt;
- Train Track;
- Tree;
- Soil;
- Standing Water.

Example feature locations are identified for each of these additional nearshore feature classes in an interactive learning session within the localized study area to expand the knowledge base support of the image mine data warehouse. A constraint in supportable feature types is that a feature class must be at least partially describable within a block of pixels of the selected snippet size. The final set of example feature labels for these 26 feature classes has about 25 000 examples (mostly generated by "area" examples selected by the sensei user, where

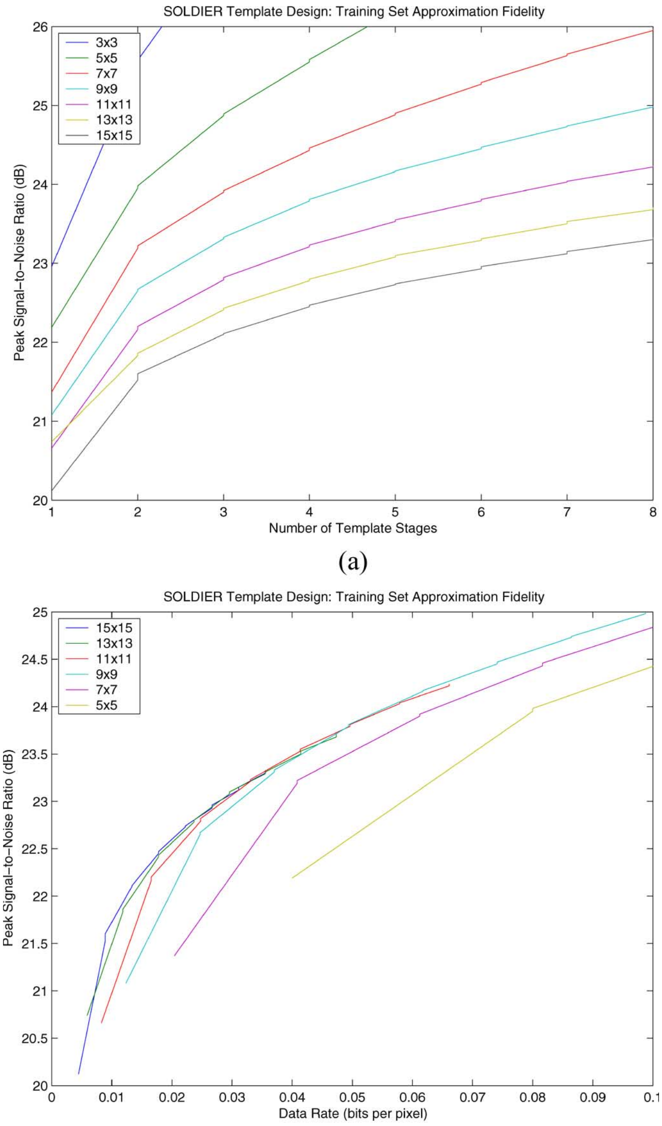

(a)

(b)

Fig. 3. Peak signal-to-noise ratios for normalized four-band pixel block sizes ranging from 3 × 3 to 15 × 15 obtained as stages are grown in the σ-tree design process. (a) SOLDIER approximation fidelity as the number of stages is increased during the design process. (b) SOLDIER approximation fidelity as the compressed data rate is increased during the design process.

each pixel in a selected homogenous feature area is entered into the feature locator table with the appropriate soft truth label). Training queries that extract all 25 000 feature example snippets of various spatial block sizes with full spectral depth (four layers) are used to generate training sets to produce σ-trees. The pixel blocks are expanded to include a negative pixel version of each spectral layer for energy normalization.

Fig. 4. SOLDIER σ-tree templates of a 26-member feature set classifier (eight-stage σ-tree with 16 templates at each stage).

TABLE I
DEBRIS FIELD CLASSIFICATION ACCURACY AND BLOCK SIZE

Block Size	Producer's Accuracy	User's Accuracy
3 × 3	86.09%	64.76%
5 × 5	81.63%	77.61%
7 × 7	81.12%	76.94%
9 × 9	76.75%	76.98%
11 × 11	62.98%	81.27%
13 × 13	60.73%	87.51%
15 × 15	56.64%	77.02%

Fig. 3 shows intermediate results of the design process as a function of block size and as the number of σ-tree layers is grown. An increasing level of aggregate approximation fidelity, measured by a peak signal-to-noise ratio, is obtained as stages are added to the σ-tree. The number of stages increases from one to eight. The design process was repeated eight times to generate eight different σ-trees as the pixel block size was varied from 3 × 3 up to 15 × 15 pixels per block. The approximation fidelity decreases as the block size is increased for a fixed number of stages. But when viewed from the perspective of data rate, measured as bits per pixel, the approximation fidelity obeys a typical rate-distortion characteristic as the rate is increased, with little difference in rate-distortion performance for block sizes of about 9 × 9 and above at a data rate of about $R = 0.04$ bits/pixel. The data rate is computed as $R = P \log_2(M)/(n \times n \times s)$, where P is the number of stages, M is the number of templates at each stage, $n \times n$ is the block spatial size in pixels, and s is the number of spectral layers. In each of these cases, the σ-tree design process produces 16 templates at each layer (see Fig. 4 for the 9 × 9 block size case). The pixel values of the residual pattern templates, as shown in Fig. 4, are real numbers and may have positive or negative values. Thus, for display purposes, these pixel templates are normalized such that the brighter pixels in the templates indicate positive values and the darker pixels indicate negative values. The residual templates at the latter stages of the template trellis have also been normalized for display purposes to show texture content.

Shown in Table I is the performance of a σ-tree detector for debris fields as a function of block size. In these experiments, the number of stages and the number of templates per stage are held fixed, and the block size is varied. The performance values in Table I show that at a block size of 9 × 9, the producer's accuracy is nearly equal to the user's accuracy (about 77%). Smaller block sizes generally produce higher producer accuracy without changing the user's accuracy. Larger block sizes generally achieve lower producer accuracy and, sometimes, higher user accuracy. This paper shows that with a σ-tree classifier that has eight stages and 16 templates at each stage, a spatial block size of 9 × 9 is a reasonable choice. All subsequent detection/classification experiment results that will be discussed in this paper for feature classes as diverse

TABLE II
CLASSIFICATION ACCURACY

Class Name	Producer's Accuracy	User's Accuracy
Residential Buildings	91%	74%
Damaged Buildings	84%	54%
Grassy Areas	90%	82%
Soil Areas	89%	99%
Roads	61%	61%
Parking Areas	76%	63%
Train Track	81%	86%
Tree	86%	70%
Road Obstructions	100%	50%
Refugee Areas	100%	100%
Sand Deposits	80%	75%
Shipping Containers	51%	82%
Standing Water	100%	68%

as buildings, roads, train tracks, sand deposits, trees, surface type, etc., were derived from the 128 templates with block size of 9 × 9, as shown in Fig. 4. All land cover/disaster impact classification accuracy estimates quantified in the remainder of this paper were evaluated with about 752 000 test snippets, none of which were used for training purposes (Table II).

C. Subfeature Fusion/Proximity Detection Filter

One of the important aspects of this paper's image mining capability is the ability to query for finely attributed features and subfeatures. The spatial snippet size used in these experiments is 9 × 9 pixels. Hence, the notion of supportable features assessable via pixel block query and the class of image object features representable by pixel block templates are those image features or subfeatures that fit within the 9 × 9 snippet block size. There are, however, many important image objects of interest that are larger than this choice of block size. A straightforward approach to recognize larger image content objects would be to increase the block size. However, when the block size is increased, computational demand increases and larger training set sizes are required for the resulting σ-trees to be well generalized. An alternative to increasing block size to deal with larger image content objects is to employ feature-level fusion and proximity analysis. That is, after the σ-tree-based maximum posterior classification of 9 × 9 features and subfeatures has been completed, a second sliding window of some size runs through the content-labeled image, and when certain combinations of subfeatures occur within the window size at sufficient frequencies, the presence of a composite higher-level object is declared. For example, the next experiment focuses on finding residential buildings with constituent subfeature fusion of shingles, building edges, and building corner subfeatures.

1) Building Detection: The subfeatures associated with residential buildings are roof edge lines, roof corners, and various residential roof colors (brown, tan, gray, dark gray shingles, etc.). Fig. 5(a) is a closeup of an area of the wide-area image scene that contains homes. This area was not used in training when the image mine warehouse was created from a localized study area. Fig. 5(b) shows the detection map of all pixel block locations that have positive detections for the following classes of subfeatures:

- Building-Roof-Edge;
- Building-Roof-Corner;
- Building-Roof-Brown;
- Building-Roof-Gray;
- Building-Roof-Dark Gray;
- Building-Roof-Tan Gray.

The detection map in Fig. 5(b) was generated by doing a σ-tree SOLDIER-MAP classification of all interior points of the image. A detection is declared, and the corresponding pixel is marked with a color if the σ-tree possesses leaf nodes that approximate the image blocks with a positive input signal energy to mismatch noise ratio [2]. The detected pixel is then assigned to the maximum posterior class over the set of 26 possible feature classes. As shown in Fig. 5(b), there is much small feature content in the image that correlates with residential building edges and corners (detections indicated with blues and purples) and with shingle texture and colors (greens and yellows). Although there is much nonbuilding-related image content that matches (both spatially and/or spectrally) these categories of subfeatures, when a subfeature fusion/proximity test is applied to retain only the detections associated with spatially adjacent subfeatures of different types related to residential buildings, the result in Fig. 5(c) is obtained. Here, detected edge or corner structures are only retained if there is a number of nearby building roof shingle detections. The proximity filter is constructed by running a sliding 9 × 9 "box window" through the image and counting the number of each subfeature detections that occur at the edges of the box. When certain combinations of subfeatures occur, e.g., shingles near building edge subfeatures, the subfeature detections are retained; otherwise, they are discarded. In these initial experiments, this is a heuristic approach without any attempt at this time to optimize performance. The sensitivity of the residential building detector is estimated to be about 91%. The specificity of the declared residential buildings is estimated to be about 74%. This classifier was not trained to detect industrial buildings, which tend to have lighter roof coverings. As seen by comparing Fig. 5(b) and (c), this proximity subfeature filter approach holds promise for false alarm mitigation and higher-level image content understanding by building on the spatial correlations of detected constituent subfeatures. An evaluation of the typical classification errors for residential buildings shows that standing water in asphalt areas is sometimes confused as residential building structures. This is believed to be a result of the spectral similarity of some roof shingle structures and wet asphalt.

2) Damaged Residential Buildings, Grass, and Soil Detection: When the fusion/proximity filter is augmented to require the presence of detected residential building roof damage (by training with a subfeature labeled as "Building-Roof-Damage") in the vicinity of roof shingles and roof edge/corner features, the wind-damaged building detection map (indicated in red) in Fig. 5(d) is obtained. The estimated wind-damage detector sensitivity is 84% with a specificity of 54%. Subjectively, debris standing in asphalt areas is sometimes confused as building damage. This is believed to be a result of the

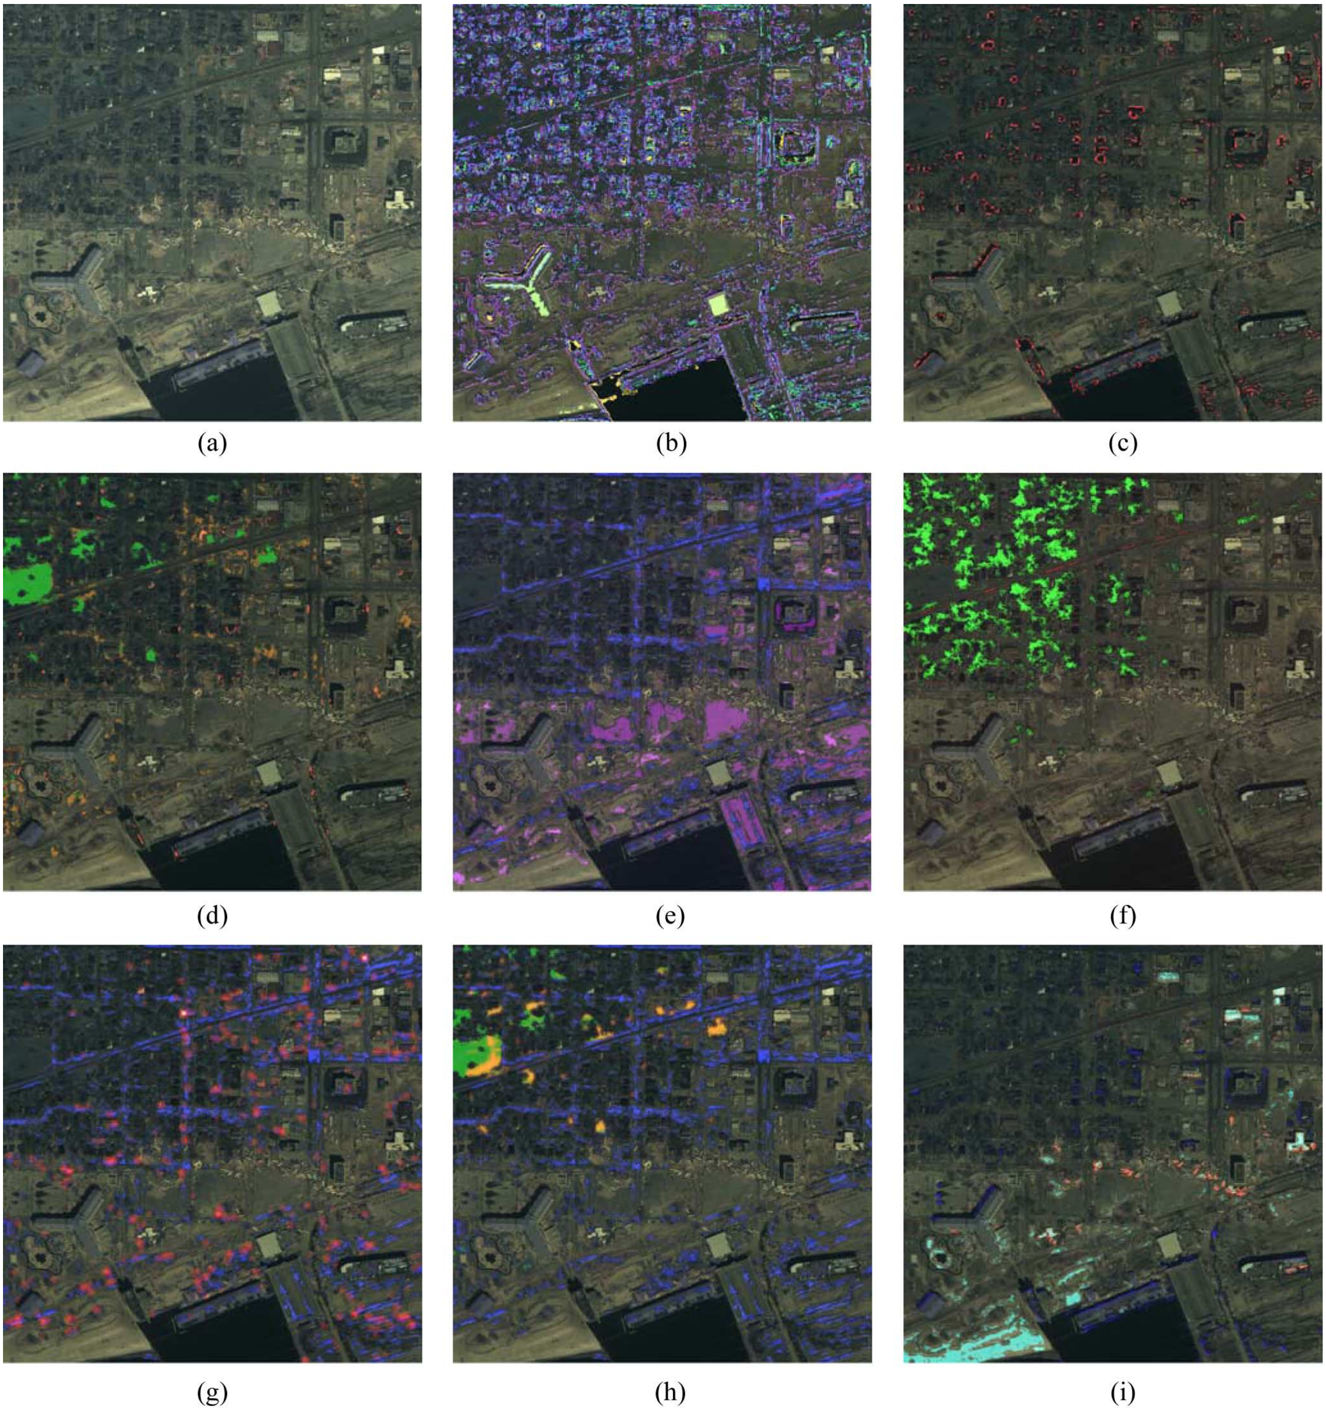

Fig. 5. Image-driven data mining results for various land cover classes. (a) Portion of large scene not included in localized training study. (b) Raw detection map of subfeatures that likely comprise residential buildings: roof shingles (green and yellow), roof edges, and corners (blue and purple). (c) House detections after proximity filtering of residential house-constituent subfeatures. (d) Wind-damaged home locations (red) determined by a spatial proximity analysis of σ-tree-detected roof shingles, roof edges, roof corners and roof damaged subfeatures, and SOLDIER-MAP detections of grassy fields (green), and bare soils (orange). (e) Raw road (asphalt and curb subfeatures) detection map shown in blue, and clear asphalt parking area detections shown in purple. (f) Train track feature detection map (red) and standing tree detection map (green). (g) Roads (blue) and obstruction points (red). (h) Raw road (asphalt and curb subfeatures) detection map shown in blue, grassy areas in green, and candidate refugee sites in orange. (i) Sand deposits (light blue) detection map, standing water (dark blue) detection map, and detections of scattered shipping containers (red).

spectral similarity of some roof shingle structures and asphalt. Also shown in this detection map are grassy areas (indicated in green) and bare soil areas (indicated in orange). The evaluated accuracy of the grass field detections is about 90% for the producer and about 82% for the user. The corresponding numbers for the soil detections are 89% and 99%.

3) Road and Asphalt Area Detection: Fig. 5(e) shows the SOLDIER-MAP detection map (blues) for subfeatures related to road structures. These features are road–asphalt and road–curb. The detection sensitivity and specificity for the combined composite class is approximately 61%. The primary confuser features for intended road detections are residential parking areas, some roof types, and some structures (parking guidelines, light debris in parking areas, etc.). A proximity filter has not been applied to the road detection result in Fig. 5(e), but it is likely that such a filter seeking to spatially correlate the road subfeatures of curb–structure, center line existence, and road surface type would do much to mitigate the apparent low specificity of this raw (no proximity filter or subfeature spatial fusion) road detection map. The figure also shows asphalt-detected parking areas (purple) with detection sensitivity of about 76% and specificity of about 63%. The primary confuser features with respect to parking areas are some colors of sand deposits and some building roof coverings.

4) Train Track and Tree Detection: Fig. 5(f) shows the detection map for the "train track" feature (red). It is difficult for an untrained observer to separate a train track from a narrow lane asphalt road, but the σ-tree image mine is able to separate these two classes of similar linear features. The specificity of this detection map is about 86% with sensitivity of 81%. There appears to be no prominent confuser class for train track detections. Visually, the train track detection is rather obvious and provides an easy visual cue to a human observer. Also shown in this figure is the detection map for trees (green). Tree detection specificity is evaluated at 70%, and the detection sensitivity is about 86%. There appears to be no prominent confuser class for hurricane-impacted tree image content.

5) Obstructed Road Detection: Fig. 5(g) shows the result of a proximity analysis performed to spatially analyze roads–asphalt, roads–curbs, and debris field subfeatures for the purpose of detecting obstructed roads. The blue features are roads and curbs, and the red features are those parts of debris fields that are sufficiently close to road subfeatures to suggest an obstruction. The obstruction detections have been slightly blurred to enlarge them and make them more visible in the detection map display. It appears that all obstructions were detected (producer's accuracy is estimated near 100%), but the specificity is only estimated to be about 50%. The false obstructions tend to be detected when the debris is near residential buildings or near wet parking areas. This is likely due to the spectral similarity of residential roof shingles and wet asphalt to asphalt road surfaces.

6) Candidate Refugee Site Detection for Rescue Tactical Support Planning: Fig. 5(h) shows the result of a proximity analysis performed to spatially cluster grass (green), roads–curbs (blue), and roads–asphalt (blue) subfeatures for the purpose of detecting candidate grassy refugee tent sites or supply staging areas with near-road access. The orange areas are those grassy (sub)areas that may be sufficiently close to road access for refugees or rescue staging activities. Although there are only a few examples for evaluation, the results show nearly 100% sensitivity and specificity for grassy areas with nearby road access.

7) Sand Deposits, Standing Water, and Scattered Shipping Container Detection: Fig. 5(i) shows areas (light blue) that seem to contain sand deposits left behind by the storm surge. Fig. 5(i) shows the raw (no proximity filter or subfeature spatial fusion) detection of sand subfeature subjectively categorized (by the sensei user) as sand deposits with tan, tan-gray, and yellow spectral signatures. The producer's accuracy is estimated to be about 80%, and the user's accuracy is evaluated at a level of about 75%. Fig. 5(i) also shows detected scattered cargo ship shipping containers (red). This small feature typically fits within the selected 9 × 9 query pixel block size. In this case, a proximity filter is not used on this feature class because the feature class has not been decomposed into more elemental subfeatures. The specificity of the shipping cargo container detections is estimated to be 82%, but the sensitivity level is estimated to be only about 51%. The low level of sensitivity is due to missed detections of the darker colored containers. If shipping containers were to be specified in terms of subfeatures with attributes such as shipping container "colors," "ends," and "centers," then a proximity filter could be applied to likely increase detection sensitivity and specificity. What few false alarms occurred tend to be near edges of industrial building structures. This figure also shows detections of standing water (dark blue) with estimated producer accuracy of nearly 100% but a user's accuracy of only 68%. Standing water detections are often confused with building shadows.

IV. CONCLUSION

Novel classification structures based on a σ-tree template trellis have been used as a foundation for systems engineering of an image-driven information mining capability for disaster assessments in high-resolution satellite imagery. The system has been demonstrated as capable of detecting a wide class of object features and subfeatures after simple end-user training exercises. Postprocessing has been demonstrated to merge detected subfeature maps for higher-level object detection and classification. This paper has shown the adaptability of the proposed σ-tree SOLDIER-MAP system to work with high-resolution satellite imagery and many fine attribute feature types without algorithm modification. The system is adaptable in that the required training phase takes only a short time and does not demand a high level of user expertise.

APPENDIX
SOLDIER BIT-PLANE INDEX GENERATION ALGORITHM

Sets of data warehouse pixel block feature exemplars $\mathbf{t}$ can be well approximated with σ-tree templates $\hat{\mathbf{t}}$. This search process results in search path index $\mathbf{i}_P$ defined as the preimage of the reconstruction mapping $\mathbf{i}_p \mapsto \hat{\mathbf{t}}$. The range of the reconstruction mapping is a set of direct sums of state templates $\{\hat{\mathbf{t}}(\mathbf{i}_P)\} = \{\hat{\mathbf{t}}_1(i_1) + \hat{\mathbf{t}}_2(i_2) + \cdots + \hat{\mathbf{t}}_P(i_P)\}$ with $\hat{\mathbf{t}}_p(\cdot) \in \Re^n$. The domain of the mapping is a set of *digital expanded representation indexes* $\mathbf{i}_P \in Z_M^P$, which is the Cartesian product of constituent M-ary *stage template index sets* $i_p \in Z_M$ with $Z_M \in [0, \ldots, M-1]$ and *stage index* $p = 1, \ldots, P$. A P-stage σ-tree search engine is indicated by $\sigma_P(\mathbf{t}) = \hat{\mathbf{t}}(\mathbf{i}_P)$, where $\mathbf{t}$ are exemplar blocks to be compressed preparatory to use in classification processing, and $\hat{\mathbf{t}}$ are directly summed approximations to the $\mathbf{t}$ of the data warehouse. The digital expanded representations $\mathbf{i}_P$ index the $\hat{\mathbf{t}}$ associated with the $\mathbf{t}$ and are produced by comparing a sequence of pixel block residuals formed from warehouse exemplar blocks $\mathbf{t}$ and the most similar stage residual templates $\hat{\mathbf{t}}_p(\cdot)$. The pth stage intermediate decision of the σ-tree sequential search engine is $\sigma_p(\mathbf{t}_p) = \hat{\mathbf{t}}(i_p)$, and the stage pattern search engine uses a local nearest neighbor decision rule, where the input to the pth *stage search engine* is the pth stage *causal residual* formed by the prior σ-tree encoder stages $\mathbf{t}_p = \mathbf{t} - \sigma_{p-1}(\mathbf{t}_{p-1}) = \mathbf{t} - \hat{\mathbf{t}}(\mathbf{i}_{p-1})$ for $p \in \{2, \ldots, P\}$ and where $\mathbf{i}_{p-1} = (i_1 \cdots i_{p-1})$ is a *partial P-tuple index*. Although a nearest neighbor rule is used locally to identify the best local σ-tree causal residual template match at each stage, there is no guarantee that the collective stage decisions $\mathbf{i}_P = (i_1, \ldots, i_P)$ provide the nearest possible $\hat{\mathbf{t}}(\mathbf{i}_P)$ for a given $\mathbf{t}$. That is, it does not necessarily follow that $d(\mathbf{t}, \hat{\mathbf{t}}(\mathbf{i}_P)) \leq d(\mathbf{t}, \hat{\mathbf{t}}(\mathbf{j}_P))$ for all $\mathbf{j}_P \in Z_M^P$. Although there is no guarantee of a global nearest neighbor search result over the set $\{\hat{\mathbf{t}}(\mathbf{i}_P)\}$ with local stage pattern search decisions, design methods presented in [1] yield σ-tree sequential search encoders with acceptable "near neighbor" performance. The tradeoff of "near" for "nearest" returns orders of magnitude reduction in implementation costs (computation and memory) [2] and is well worth a degree of global suboptimality. Furthermore, whatever measure of joint suboptimality is incurred in this tradeoff can be compensated for by increasing the number of stages, which requires only a linear complexity expenditure in both computation and memory.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for suggestions that have improved this paper's content and scope.

REFERENCES

[1] C. F. Barnes, S. A. Rizvi, and N. M. Nasrabadi, "Advances in residual vector quantization: A review," *IEEE Trans. Image Process.*, vol. 5, no. 2, pp. 226–262, Feb. 1996.

[2] C. F. Barnes, "Image driven data mining for image content segmentation, classification and attribution," *IEEE Trans. Geosci. Remote Sens.*, 2007, to be published.

[3] C. Chang, B. Moon, A. Acharya, C. Shock, A. Sussman, and J. Saltz, "Titan: A high-performance remote sensing database," in *Proc. 13th Int. Conf. Data Eng.*, 1997, pp. 375–384.

[4] R. G. Congalton and K. Green, *Assessing the Accuracy of Remotely Sensed Data: Principles and Practices*. Boca Raton, FL: CRC, 1999.

[5] M. Datcu, H. Dashiel, A. Pelizzari, M. Quartulli, A. Galoppo, A. Colapicchioni, M. Pastori, K. Seidel, P. G. Marchetti, and S. D'Elia, "Information mining in remote sensing image archives: System concepts," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 12, pp. 2923–2936, Dec. 2003.

[6] U. M. Fayyad, G. P. Shapiro, P. Smyth, and R. Uthurusamy, *Advances in Knowledge Discovery and Data Mining*. Cambridge, MA: AAI/MIT Press, 1996.

[7] J. Li and R. Narayanan, "Integrated spectral and spatial information mining in remote sensing imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 3, pp. 673–685, Mar. 2004.

[8] G. B. Marchisio and J. Cornelison, "Content-based search and clustering of remote sensing imagery," in *Proc. IGARSS*, Hamburg, Germany, Jun. 1999, vol. 1, pp. 290–292.

[9] C. R. Veltkamp, H. Burkhardt, and H.-P. Kriegel, *State-of-the Art in Content-Based Image and Image Retrieval*. Norwell, MA: Kluwer, 2001.

[10] J. Zhang, H. Wynne, and M. L. Lee, "Image mining: Issues, frameworks and techniques," in *Proc. 2nd Int. Workshop Multimedia Data Mining*, San Francisco, CA, Aug. 2001, pp. 13–20.

Christopher F. Barnes (S'85–M'87–SM'91) received the Ph.D. degree in electrical engineering from Brigham Young University, Provo, UT, in 1989.

In 1989, he joined the Georgia Tech Research Institute (GTRI), Atlanta, and worked there for nearly 13 years. He transferred within Georgia Tech from GTRI to Georgia Tech's Regional Engineering Program (GTREP), Savannah, as an Associate Professor in the summer of 2002 and now holds a joint appointment with Georgia Tech's School of Electrical and Computer Engineering. He has over 17 years experience in basic and applied research. His research interests include information and communication theory, signal processing, and machine learning. His areas of expertise include data compression, trellis-coded modulation, synthetic aperture imaging, pattern recognition, and high-performance computing for ground-based radar signal and data processing.

Hermann Fritz received the Dr.Sc.Tech. (Ph.D. equivalent) degree from the Swiss Federal Institute of Technology, Zurich, in 2002.

In March 2003, he started his dual appointment as an Assistant Professor with the Georgia Tech Regional Engineering Program, Savannah, and the School of Civil and Environmental Engineering, Atlanta, GA. His research centers on hydrodynamic aspects of natural hazards such as tsunamis, landslides, and hurricanes. He has led and participated in numerous International Tsunami Survey Teams in the aftermath of the 2004 Indian Ocean Mega-Tsunami. He has surveyed the tsunami impact in Sumatra (Indonesia), Sri Lanka, The Maldives, Somalia, Madagascar, and Oman and led the largest independent and nongovernmental Hurricane Katrina storm surge survey campaign along the Louisiana, Mississippi, Alabama, and Florida Gulf shores.

Jeseon Yoo received the B.S. degree in oceanography from Seoul National University, Seoul, Korea, in 1996 and the M.S. degree in civil and environmental engineering from Korea Advanced Institute of Science and Technology, Daejon, Korea, in 2002. He is currently working toward the Ph.D. degree in civil and environmental engineering at the Georgia Institute of Technology, Savannah.

His research interests are in estimating water wave property and modeling water depth inversion by using linear feature extraction from nearshore wave images.