

AUTOMATIC SINKHOLE DETECTION FROM SATELLITE- BASED INSAR

A Thesis
Presented to
the faculty of the School of Engineering and Applied Science
University of Virginia

In Partial Fulfillment
of the requirements for the Degree
Master of Science in Electrical Engineering

By

Michael Stuecheli
May 2013

Approval Sheet

This thesis is submitted in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering

Michael Stuecheli – Author

This thesis has been read and approved by the examining Committee:

Scott T. Acton – Thesis Advisor

Harry C. Powell – Committee Chair

Gang Tao – Committee Member

Accepted for the School of Engineering and Applied Science:

Dean, School of Engineering and
Applied Science

May 2013

Acknowledgements

I would like to extend my sincere thanks to my advisor, Professor Acton, for his support and guidance during my long stay at this university. Under his tutelage, I have garnered a decent understanding of the broad field of image processing, and feel confident that I can carry on learning and applying this knowledge in future endeavors. I am also grateful to the various members of my lab, VIVA—especially the future Dr. Dre, Andrea. The work contained herein is a product of the arduous efforts of myself, Andrea, and of course Dr. Acton. Lastly, I would like to thank the members of my committee, Dr. Tao and Dr. Powell, both for overseeing this final step in my education as well as their efforts in the many classes and projects which I participated in and enjoyed during my time at the University of Virginia.

Abstract

Recent developments in both the quality and processing of Interferometric Synthetic Aperture Radar (InSAR) data have allowed for remote analysis of ground deformation with unprecedented accuracy. SAR images acquired via the COSMO-SkyMed satellite constellation provide spatial resolutions on the order of a couple of meters per pixel; with a reasonable number of these images, SAR-processing algorithms known as PSInSAR and SqueeSAR can yield ground deformation information with millimeter accuracy. In this thesis, I present two novel approaches which utilize the aforementioned ground displacement measurements in order to detect and monitor geological hazards known colloquially as sinkholes. The first algorithm employs a graph theoretic approach to accomplish this goal, while the second utilizes a spatiotemporal parametric matching approach; both of these methods demonstrate strong efficacy in locating sinkholes from satellite-based InSAR-derived data.

Contents

Acknowledgements	iii
Abstract	iv
List of Figures	vi
List of Tables	viii
1. Introduction.....	1
1.1 InSAR-based General Detection of Geohazards.....	1
2. Automatic Sinkhole Detection via Graph Cuts.....	4
2.1 Graph Cuts.....	4
2.2 Method	5
2.3 Results and Analysis	10
2.3.1 Wink Sinks.....	10
3. Automatic Sinkhole Detection via Parametric Spatiotemporal Matching.....	14
3.1 Sinkhole Model	14
3.2 Method	19
3.3 Results and Analysis	22
3.3.1 Validation—Simulated Data.....	22
3.3.2 Wink Sinks.....	26
4. Approach Comparison	30
4.1 Simulated Data	30
4.2 Wink Dataset.....	33
4.3 Virginia Data—General Comparisons	35
5. Conclusions.....	39
5.1 Applicability/Generality.....	40
Bibliography	41

List of Figures

Figure 2-1 The graph cut segmentation process, pictorially. The input displacement point cloud (a) is transformed into a weighted graph structure (b), the optimal cuts located (c) and final segmentation overlayed (d).....	8
Figure 2-2: Aerial imagery of the area imaged for the Wink dataset	10
Figure 2-3: Temporal, graph-cut segmentation routine processing on the Wink dataset; (a) displays the global score, (b) the temporal score, (c) the combination of the two, and (d) the overlay of the binned results	12
Figure 3-1: Example profile generated from Wink 3 at a sampled time t ; the same general color scale applies as in Figure 2-1.....	14
Figure 3-2: Profile generation for Winks 2 and 3	16
Figure 3-3: Normalized cross-correlations between profiles and their respective Gaussian fits	17
Figure 3-4: Temporal evolution of amplitude (a) and sigma (b)	18
Figure 3-5: Temporal evolution of amplitude vs. linear regression.....	18
Figure 3-6: Simulated Datasets	22
Figure 3-7: Minimum Residual Images	24
Figure 3-8: Propagated minimum residual images	26
Figure 3-9: Propagated minimum residual image for the Wink dataset (a), corresponding propagated maximum risk image (b)	28
Figure 3-10: Binned propagated maximum risk overlaid upon visual satellite imagery..	28
Figure 4-1: Simulated datasets.....	30
Figure 4-2: Comparison of the two methods on simulated dataset.....	32
Figure 4-3: Resulting segmentation of the Wink dataset for the graph cut algorithm (a) and the spatiotemporal matching approach (b).	34
Figure 4-4: Detailed view of Winks 1 and 2 (top and bottom rows) for graph cut algorithm (a,c) and spatiotemporal matching (b,d).....	35

Figure 4-5: Example overlap and uniqueness between graph cut detector (a) and matching algorithm (b)	38
--	----

List of Tables

Table 4-1: Overall detection distributions	36
Table 4-2: Detection overlap between methods.....	37

1. Introduction

Headlines describing an anthropomorphic earth opening up and swallowing people are quite the attention-grabber, as are the geological hazards which cause them—sinkholes. The term ‘sinkhole’ refers to a ground depression which occurs in Karst geologies; Karst areas are composed of carbonate rocks such as limestone, dolomite, etc. The action of slightly acidic water leads to the dissolution of these underlying rocks, which in turn can eventually lead to a sudden or gradual formation of a sinkhole. These phenomena can deleteriously impact housing, transportation, and various infrastructures—sometimes even in a life-threatening way. Importantly, it is possible to locate these sinkholes before they collapse, and remediate them accordingly [1]. In this thesis, I describe two automatic approaches which aim to accomplish this goal.

1.1 InSAR-based General Detection of Geohazards

There are numerous potential routes which can be taken in an effort to detect sinkholes; while the foremost amongst these—electrical resistivity, ground penetrating radar, microgravimetry, and cross-hole tomography [2]—are certainly viable options in estimating subsurface soil/rock/void characteristics and locations, they present salient disadvantages in terms of speed, efficacy, and coverage. However, another method is available in the pursuit of sinkhole location and monitoring: Interferometric Synthetic Aperture Radar, InSAR. In particular, satellite-based InSAR is well-suited for this application, as it is capable of nearly instantaneously surveying large areas with a reliable repeat cycle. This is accomplished by acquiring at least three SAR images of the same

scene at different times, then calculating the phase difference between all combinations of these images—once these resulting images have been processed (phase unwrapping, atmospheric corrections, etc) it is possible to determine the elapsed vertical displacement of the ground with respect to the line of sight of the satellite to an accuracy of centimeters [3].

In this study, we employ SAR data which has been processed much more extensively than simply combinatorially evaluating SAR images. With the application of an algorithm developed by a research partner, TRE, known as PSInSAR/SqueeSAR, we are able to utilize displacement data which provides information to an accuracy of millimeters. In short, these proprietary algorithms process large stacks of SAR images (on the order of 20+) of a scene, identify permanent scatterers in the scene, and use these to correct various distortions such as atmospheric variations [4]. Additionally, the algorithm also identifies statistically homogenous regions which exhibit similar scattering characteristics (distributed scatterers) and is able to extract accurate displacement characteristics for that area [5]. The result of this processing is a point cloud of spatiotemporal displacement data where each spatial point corresponds to either a PS (permanent scatterer) or DS (distributed scatterer); moreover, this point cloud is both sparse and non-uniformly distributed. This is the finalized data upon which our algorithms operate.

Hazard detection via satellite-based InSAR has seen some exposure in the literature; there are a couple of previous studies which have taken this approach. Paine et. alia [6] utilized InSAR in order to find subsiding locations of interest, then deployed microgravity measurements to further characterize those locations. The addition of

microgravimetry to finely characterize potential sinkholes is distinct from our study, where we defer to more generalized field verification carried out by a geologist. However, the process of acquiring these microgravity measurements is painstaking and requires direct access to a site—i.e. one must physically drive over the location, after setting up calibration stations on either side of the measurement path. Additionally, the aforementioned study only uses a single SAR interferogram upon which potential detections are identified; as such, this interferogram provides rough displacements on the order of centimeters. Our study employs a much more sophisticated process in addition to a large stack of SAR images, and we operate upon a time-series of displacement information accurate to the order of millimeters.

Another group [7] employed InSAR in Arizona in order to measure subsidence at various sites of interest. Their study used three InSAR images acquired over three years, and after applying phase-unwrapping, DEM correction, and speckle reduction techniques, claimed an ability to measure displacements on the order of millimeters. The authors identified four potential subsidence features with their method. However, this method does not make any efforts to numerically characterize any of these features, nor does it do so automatically. Additionally, the resolution of the SAR images used is on the order of 30m x 30m, and combined with their application of a spatial moving average filter in smoothing noise, yields a very low amount of spatial information. This is in stark contrast to our SAR images and respective processing, which yield a pixel resolution of a couple of meters.

2. Automatic Sinkhole Detection via Graph Cuts

In this section, I present the first of two potential avenues for the automatic detection of sinkholes; in both cases, this detection is carried out upon InSAR image stacks processed with the aforementioned PSInSAR and SqueeSAR algorithms. Thus, the following approach deals with the sparse, nonuniformly-sampled, spatio-temporal displacement point clouds.

2.1 Graph Cuts

The sparse and nonuniform nature of our data precludes the possibility of directly employing typical image processing techniques—however, approached from a graph-theoretic perspective, these Cartesian aspects of the data become irrelevant. One can efficiently and meaningfully process as well as make conclusions with the rich set of tools which graph theory provides.

Specifically, we employ a powerful subset of graph theory known as the graph cut. A graph cut is a segmentation of an arbitrary graph via the ‘cutting’ of edges; the optimal cut is defined as that which yields minimum energy

$$E_{cut} = \sum_{u \in A, v \in B} w(u, v) \quad (2.1)$$

where A and B are the disjoint sets of vertices created by the cut, and $w(u, v)$ is the weight of a given edge between vertices (u, v) . Typically, the energy of a cut is expanded to include a measure of association between the ‘cut’ sets [8] A and B —this promotes spatial coherence, preventing small, isolated vertices from being unnecessarily cut away. This expansion is sometimes termed the normalized cut:

$$NE_{cut} = \frac{E_{cut}(A, B)}{assoc(A, V)} + \frac{E_{cut}(A, B)}{assoc(B, V)} \quad (2.2)$$

and

$$assoc(A, V) = \sum_{u \in A, z \in V} w(u, z) \quad (2.3)$$

$$assoc(B, V) = \sum_{u \in B, z \in V} w(u, z) \quad (2.4)$$

In (2.3) and (2.4), z is any other vertex in the set of all vertices, V .

Alternatively, the min-cut partition can be found via its strong optimization dual, maximizing flow in the graph (max-flow). Flow can most intuitively be grasped in terms of a water analogy: edges function as pipes for flow, with capacities for flow equal to their respective weights. Predetermined nodes called sources produce flow, while nodes termed sinks absorb flow. With the aid of a clever labeling algorithm, calculating the state which maximizes the flow from the sources to the sinks also identifies the minimum-energy cuts in the graph [9]—the sources and sinks reside in disjoint sets A and B , respectively. Thus an equivalent expression to (2.2) is the functional [10]

$$E_{cut}(f) = \lambda \sum_{p \in P} R_p(L_p) + \sum_{\{p, q\} \in N} B_{\{p, q\}}(L_p, L_q) \quad (2.5)$$

where p and q are vertices, B is an associativity term, R is analogous to the traditional energy assigned to an edge, and λ is a weighting factor.

2.2 Method

In order to leverage graph cuts in the automatic analysis of our data, unsurprisingly one must first construct a graph. Doing so with a sparse point cloud is straightforward; the points are taken as vertices, and edges are derived via a Delaunay triangulation. The

two-dimensional Delaunay triangulation has the virtue of maximizing the minimum angle of each triangle produced—this avoids the creation of long ‘skinny’ triangles—in addition to creating non-overlapping edges. These two features result in an edge construct which accurately and densely reflects local neighborhood relationships within the point cloud.

Segmentation via graph cuts is by no means a novel development; the approach is very well-represented in the literature. However, this segmentation is typically applied to the realm of images—the graph is constructed from rasterized data, and segmentation is based upon pixel intensity and location information. This presents a large informational disparity with respect to our unique data, which supports a much broader feature set. In addition to location information, each point also has an associated displacement time series, coherence, velocity, acceleration, physical area represented, etc. This amounts to a greatly expanded realm of possibilities for edge weighting schemes. As a result, it became possible to devise a clever edge-weighting functional which allows for the detection of spot-like, subsiding regions—i.e. sinkholes. In order to accomplish this, the energy functional must favor a minimum when on an edge which intersects the boundary of a spot-like, subsiding region. Thus, minimum-energy cuts will follow these boundaries and correspondingly segment suspected sinkholes in the desired fashion.

Three major nodal properties were chosen to accurately indicate the presence of a sinkhole-forming region:

- displacement difference between two nodes
- displacement range
- coherence

An edge which resides on the boundary of a spot will have a large displacement difference between its respective nodes, and will encompass a specified negative displacement range—essentially a level set. The coherence is a direct measure of the accuracy of a node, and thus helps to ensure that the most accurate points drive the detections (or lack thereof). For each edge, e_i , connecting vertices (v_{i1}, v_{i2}) with displacements (d_{i1}, d_{i2}) where $d_{i2} > d_{i1}$, and average coherence coh_{avg} , we arrive at:

$$E_{band}(v_{i1}, v_{i2}) = \begin{cases} 0 & \text{if } d_{i1} < d_c < d_{i2} \\ \min(|d_{i2} - d_c|, |d_{i1} - d_c|) & \text{else} \end{cases} \quad (2.6)$$

$$E_{diff} = \frac{1}{|d_{i2} - d_{i1}|} \quad (2.7)$$

$$E_{reliable} = \frac{1}{coh_{avg}} \quad (2.8)$$

$$E_{length} = length(e_i) \quad (2.9)$$

E_{length} is included to give priority to cuts that pass through shorter, and thus more accurate, edges, in turn increasing the accuracy of the cut. In order to facilitate a meaningful combination of these energies, each energy is normalized and combined to form a total edge energy

$$E = \alpha \cdot E_{band} + \beta \cdot E_{diff} + \gamma \cdot E_{reliable} + \delta \cdot E_{length} \quad (2.10)$$

where α , β , γ , and δ are inter-category weighting parameters. This energy intentionally omits an associativity measure, as spatial coherence is not a concern in this application—individual points can potentially represent quite large areas, and potential detections which include only these points are also factoring in numerous neighborhood relations.

Once this conglomerate energy function has been calculated, the minimum-energy graph cuts are then located. After fitting a spline to each cut to create smooth contours, the final segmentation of spot-like, subsiding regions is achieved. This entire process is

summarized in Figure 2-1, where $\alpha = 0.1$, $\beta = 1$, $\gamma = 2$, $\delta = 0.5$, and $d_c = -107$ mm. An input displacement point cloud at time t is shown in (a), the graph structure and resulting edge weights are shown in (b), where the color scale retains the same proportional meaning. The segmentation of the subsiding spot-like region is then overlaid upon the graph structure (c) and the original point cloud (d).

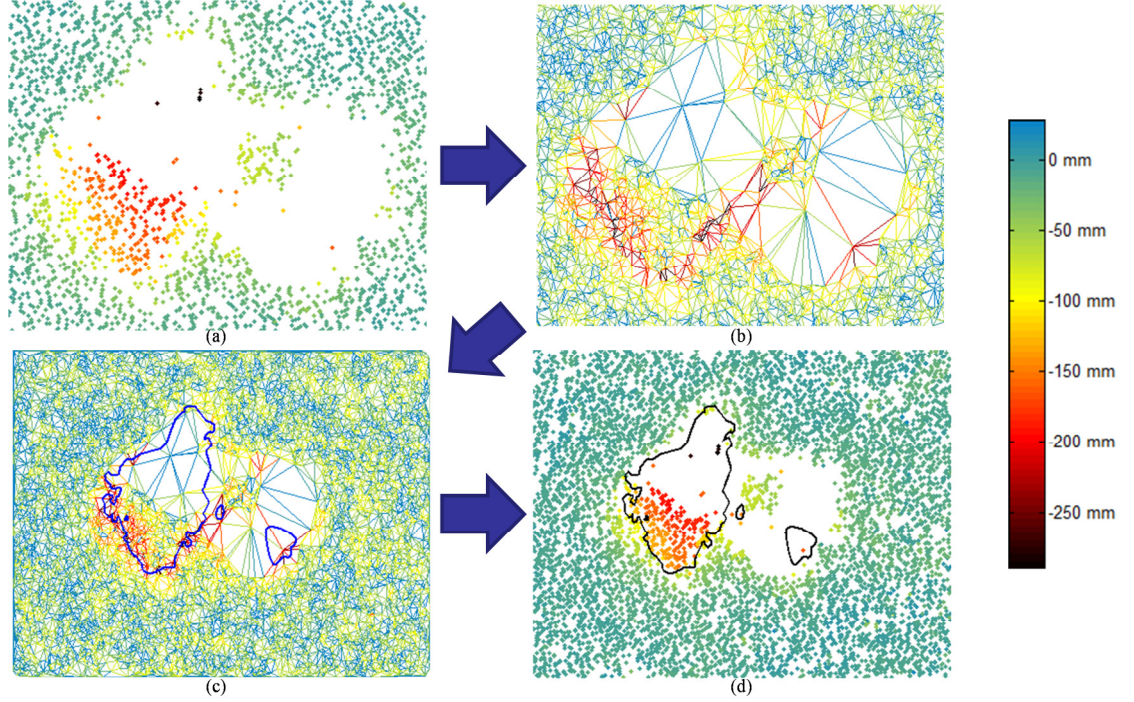


Figure 2-1 The graph cut segmentation process, pictorially. The input displacement point cloud (a) is transformed into a weighted graph structure (b), the optimal cuts located (c) and final segmentation overlaid (d)

At this point, the algorithm is capable of detecting spot-like, subsiding regions at a specified time t and spot depth cutoff d_c . In order to expand this analysis to include multiple significant displacement depths as well as temporal evolution of these suspected sinkholes, the sparse spot detection described above is implemented in a highly iterative fashion and its outputs are combinatorially analyzed.

To this end, the algorithm seeks to capture both the global—in a time sense of the word—behavior of the point cloud in addition to its more mercurial aspects. In measuring these characteristics, one must still account for the level-set ‘slice’ aspect of the sparse

spot detection. Intuitively, subsiding spot-like regions which are consistently detected across numerous slice level sets present a more coherent threat; additionally, those regions which also present a detection for lower level sets are further increased in precedence. Thus, for a given temporal sample of interest, one can iteratively run the previously described sparse spot detector for a range of displacement level sets, and assign a proportionally higher detection value to those regions which are segmented across multiple levels—with an emphasis upon deeper levels.

The algorithm proceeds as follows:

- Measure temporally global behavior by analyzing the final displacement point cloud—normalize these displacements by time elapsed, then iteratively employ the spot-like subsiding region detector with a range of (normalized) displacement level sets $d_{c1}, d_{c2}, \dots, d_{cN}$. Lastly, overlap detected regions and assign a detection value proportional to the number of overlaps which occur, as well as the deepest d_c at which a region was first detected.
- Measure temporal evolution by analyzing time-windowed, normalized displacements. For a dataset with M displacement time entries, employ a sliding temporal window which captures three contiguous displacement time entries. For each spatial point, calculate a normalized displacement via a linear regression in time. Next, iteratively employ the sparse spot detector (as in the global case) with a range of normalized displacement level sets $d_{c1}, d_{c2}, \dots, d_{cN}$, and assign detection values for each window as described in the case above.
- Combine the detection scores for the $M - 3$ windows via a simple average, then synthesize a final score with the combined window scores and the global score

$S_{final} = \varepsilon \cdot S_{global} + \mu \cdot S_{window}$, where ε and μ determine the proportional influence of each factor; this could depend upon the types of sinkholes being looked for, whether one wishes to include the influence of other temporal facets such as rainfall, etc.

2.3 Results and Analysis

2.3.1 Wink Sinks

The aforementioned algorithm was tested upon a training dataset supplied by TRE Canada known as the Wink Sink dataset.

This data is derived from 22 InSAR images acquired by the ERS satellites from 1992 through 1998; interestingly, the area imaged contains four notable geological hazards. The first hazard, Wink 1, is the large sinkhole at top in Figure 2-2 which collapsed in 1980 (before the data was acquired). The second, Wink 2, is also a massive sinkhole which had the virtue of collapsing after the acquisition period, in 2002. This

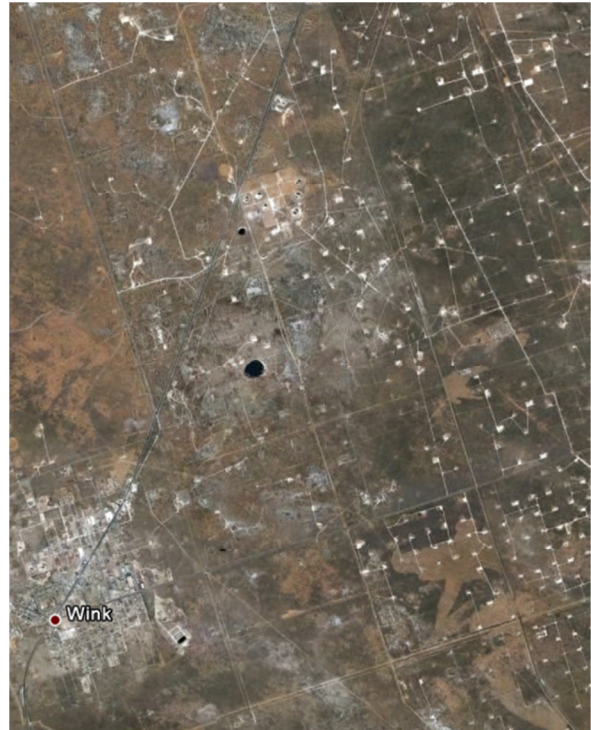


Figure 2-2: Aerial imagery of the area imaged for the Wink dataset

sinkhole is the larger black spot in the middle of the image. The last two hazards, Winks 3 and 4, do not present a large footprint in the visible light spectrum, but are strongly suspected to be sinkholes in a pre-collapse state. This 22-image stack was processed with

PSInSAR and SqueeSAR analysis, yielding 95,513 data points over an area 55.37 km² in size.

Parsing this data cloud with the graph-cut-based segmentation algorithm yielded the sequence of images on the following page. Parameters implemented were as follows: $\alpha = 0.1$, $\beta = 1$, $\gamma = 2$, $\delta = 0.5$, $d_{c1} = -4\text{mm/month}$, $d_{c2} = -3.75\text{ mm/month}$, \dots , $d_{c15} = -0.5\text{ mm/month}$, $\varepsilon = 0.25$, and $\mu = 0.75$. Figure 2-3a illustrates the detection score from the global analysis, S_{global} ; scored regions are here depicted by the data points which lie in their interiors, and a score of 0 represents the most severe potential sinkhole detection. The global scoring picks up the collapsed sinkholes Wink 1 and Wink 2 well, and also is able to detect the suspected sinkholes Winks 3 and 4. Figure 2-3b shows the detection score from the temporal analysis, S_{window} , which seems to support the results from the global analysis. However, this temporal version picks up far more of a response from other areas which perhaps exhibit behaviors associated with sinkholes but in a less coherent, significant way. Combined, these two corroborate the original main four detections (Figure 2-3c) while still giving some credence to other potential movements.

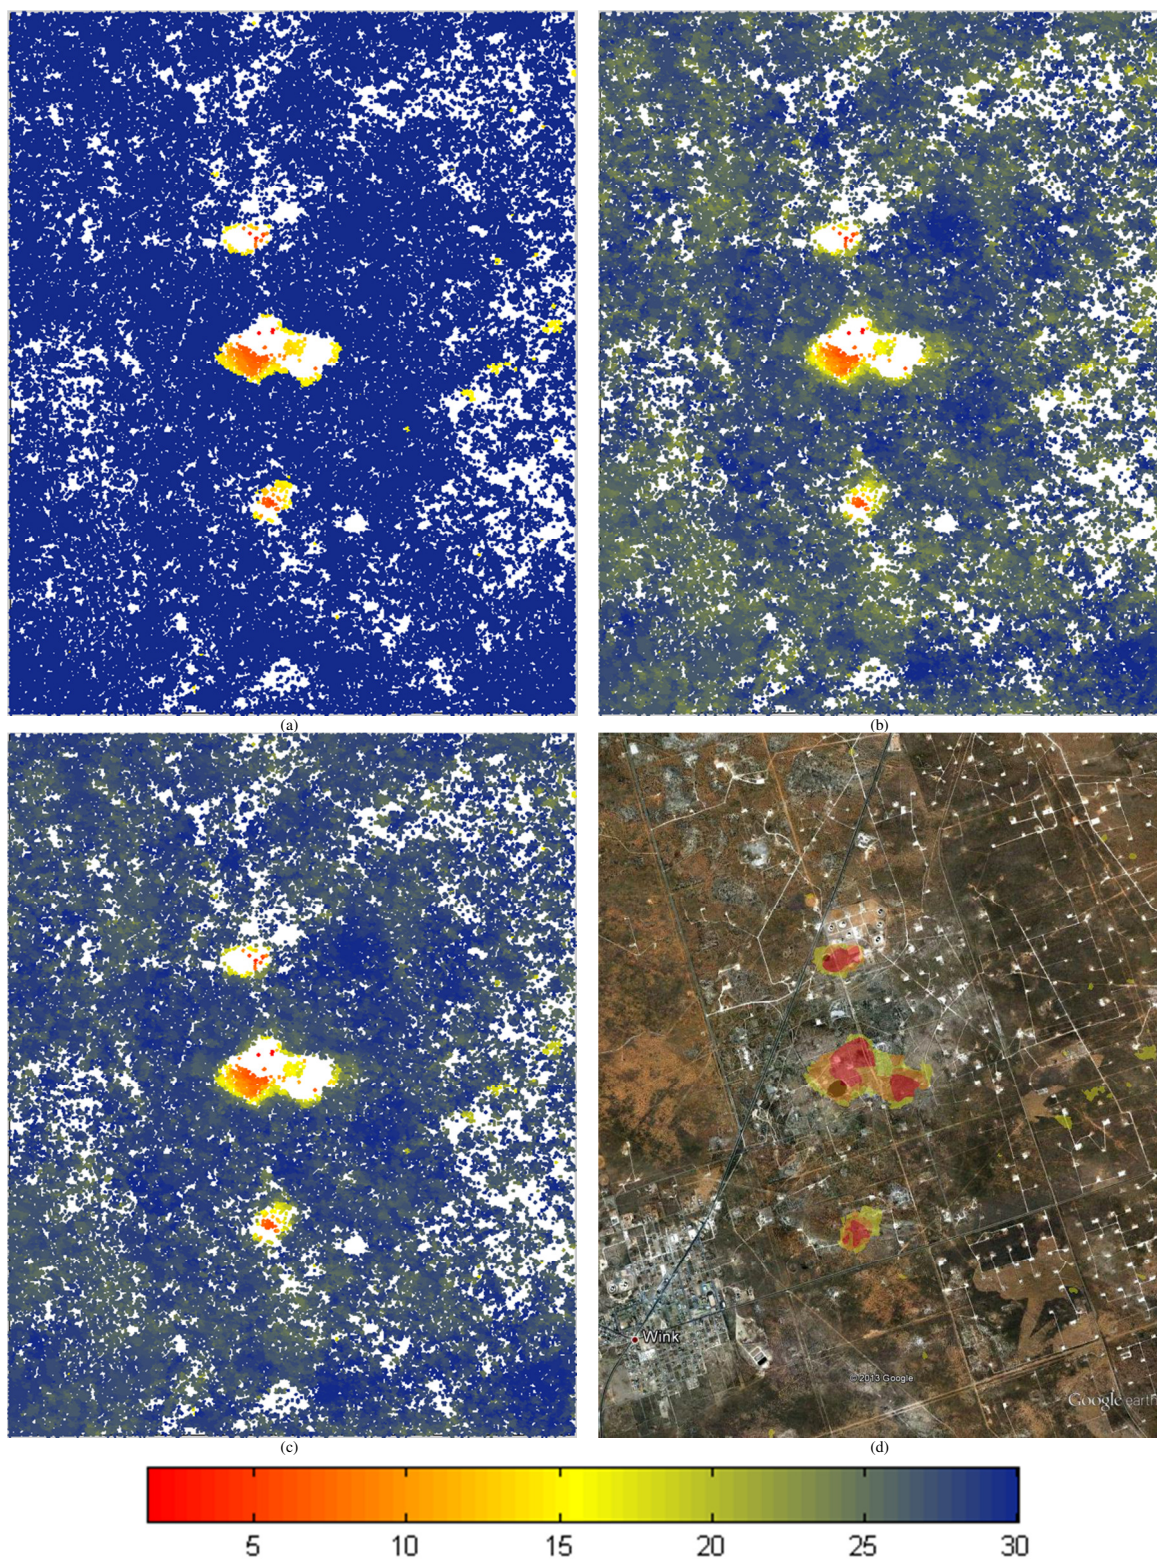


Figure 2-3: Temporal, graph-cut segmentation routine processing on the Wink dataset; (a) displays the global score, (b) the temporal score, (c) the combination of the two, and (d) the overlay of the binned results

In the last image, Figure 2-3d, the combined detection scores are binned into four main detection rankings, and overlaid upon satellite imagery. Here red corresponds to severe risk, orange with moderate risk, and yellow with slight risk. The fourth risk color, blue, is completely transparent in this depiction. With this binning, the algorithm pinpoints all four of the geohazards as mostly severe-moderate detections, and exhibits some minor detections on the periphery of the image as well. Specifically, 100% of Wink 1 is detected—37.5% is labeled as severe, 57.8% as moderate, and 4.7% as slight; 100% of Wink 2 is detected—17.2% is labeled as severe and 82.8% is labeled as moderate. It is worth noting here that Wink 1, because it collapsed before the dataset was acquired, presented very little in the way of coherent radar scattering characteristics, and thus yields few points in the vicinity of the actual hole. This most likely accounts for the increased difficulty in capturing a greater portion of the hole with the severe label.

3. Automatic Sinkhole Detection via Parametric Spatiotemporal Matching

In this next section, I discuss the second algorithm for the automatic detection of sinkholes based upon spatiotemporal point cloud displacement data derived from PSInSAR/SqueeSAR processing of InSAR image stacks.

3.1 Sinkhole Model

In an effort to generate a formal spatiotemporal model for a solution sinkhole, it became advantageous to generate and analyze profiles of such known sinkholes. To this end, the Wink Sink dataset was briefly employed in an observational, exploratory analysis.

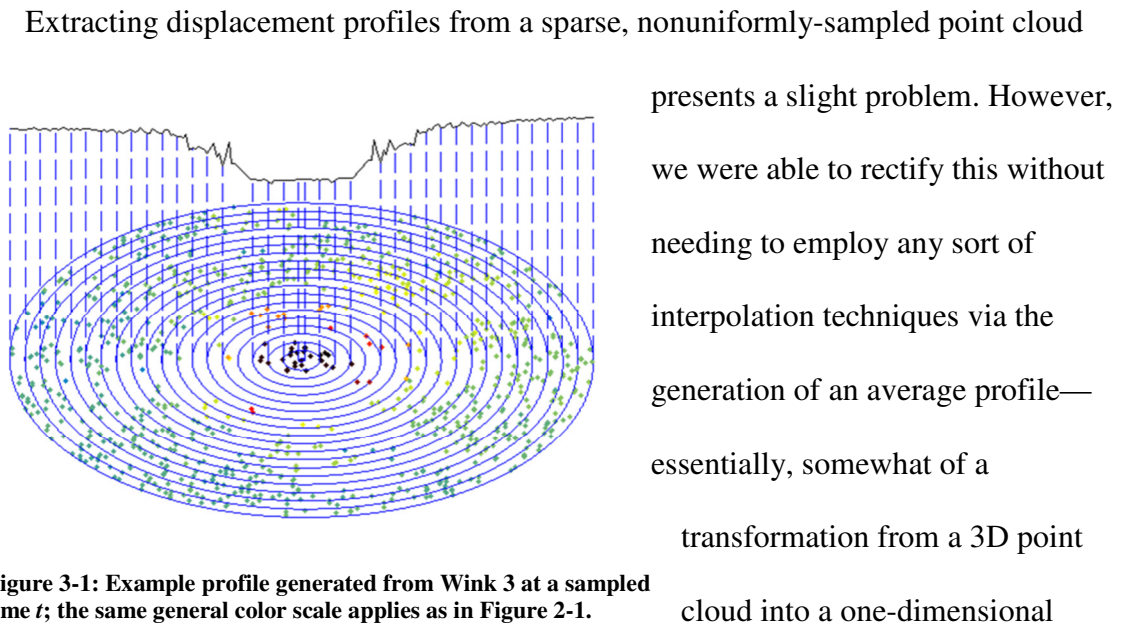


Figure 3-1: Example profile generated from Wink 3 at a sampled time t ; the same general color scale applies as in Figure 2-1.

profile. This was accomplished by generating a series of concentric annuli, centered approximately on the origin of a known sinkhole. Each annulus is about five meters in

width, and any points falling within are simply averaged together to form a new displacement point on the 1D displacement profile (Figure 3-1).

This profile generation was carried out for Wink Sinks 1-3, and for each available time sample. The spatial extent of the profile generation reaches somewhat far beyond the actual area covered by the sinkholes (and suspected sinkhole), as each geological hazard wields influence significantly beyond its collapsed region. Figure 3-2 displays the results for Winks 2 and 3, which are most densely graced with data points. In (a) we see the spatiotemporal plot of these profiles for each Wink, where every solid line in the spatial axis corresponds to a temporal sample, and every solid line in the time axis corresponds to an annulus location. Upon projecting to the spatial axis, it becomes evident that these profiles exhibit a strongly Gaussian shape throughout their lifetimes (c), and projecting to the temporal axis interestingly seems to suggest a definite linear behavior in the evolution of these sinkhole profiles (b).

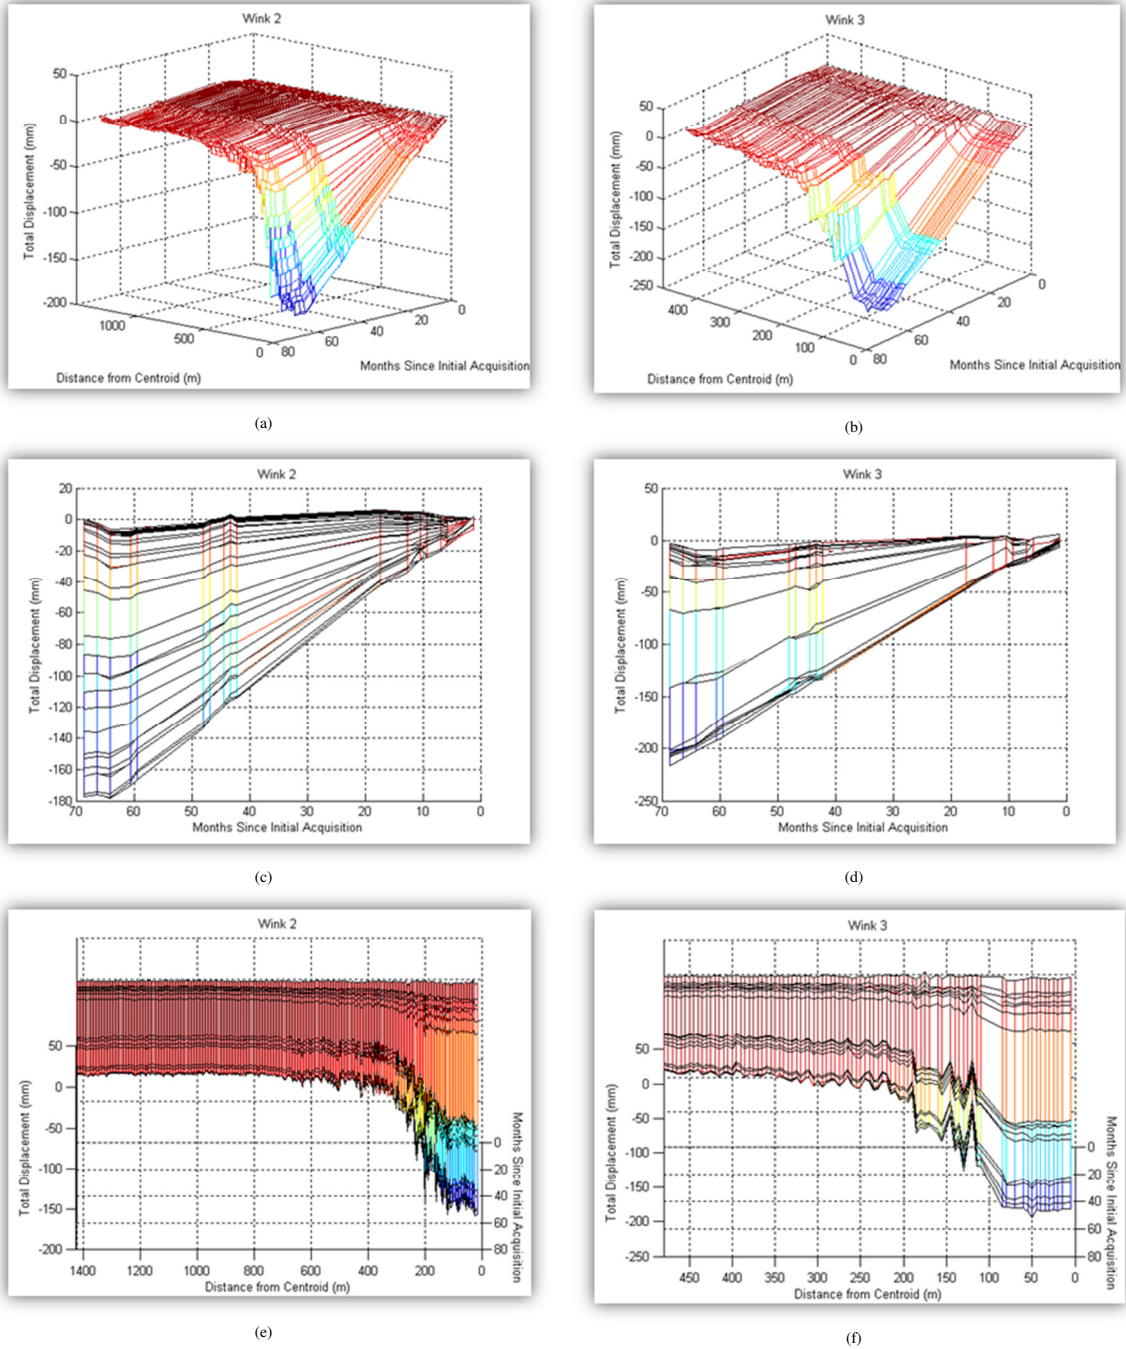


Figure 3-2: Profile generation for Winks 2 and 3

In pursuit of demonstrating the validity of the exhibition of spatially Gaussian profiles in displacement trends of forming sinkholes, we first attempt to fit our generated profiles (above) with Gaussian curves of amplitude α_t and standard deviation (scale parameter) σ_t .

$$g_t(x) = \alpha_t \exp\left(\frac{-x^2}{2\sigma_t^2}\right) \quad (3.1)$$

Here the subscript t indicates that these parameters were evaluated independently at each time t .

With these fits in hand, we then numerically substantiate the suitability of this Gaussian approximation by calculating the normalized cross-correlation between the actual profile and fit Gaussian at each timeframe.

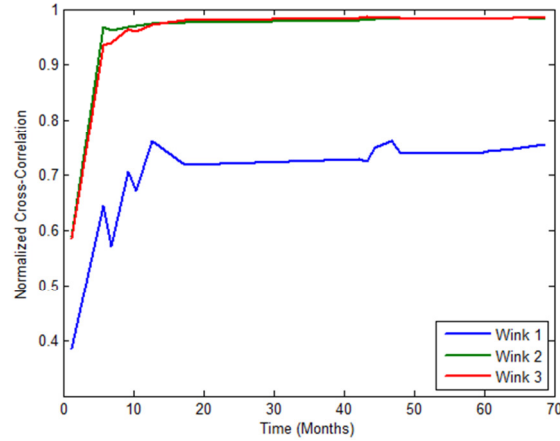


Figure 3-3: Normalized cross-correlations between profiles and their respective Gaussian fits

The results are shown for Winks 1-3

(Figure 3-3). In all cases, the correlation coefficient was relatively weak in the first few time frames ($t < 7$ months), as at that point there is no significant deformation which has yet occurred and the measurements are contending with noise. However, once beyond this point in time, Winks 2 and 3—both of which exhibit a strong sampling of data points in their immediate and surrounding vicinities—exhibited correlation coefficients approaching unity. Wink 1 on the other hand, as discussed earlier, collapsed before the dataset was acquired; the hole presented very poor scattering characteristics for our acquisitions as it had, among other things, filled with water. Thus we see a lesser correlation at about 0.75—this is still encouraging, as the points surrounding Wink 1 still indicate a Gaussian shape despite the lack of any data within about 1 - 2 standard deviations of its optimal model fit.

Next, we analyzed the temporal evolution of the amplitude and standard deviation parameters for each Wink Sink (Figure 3-4).

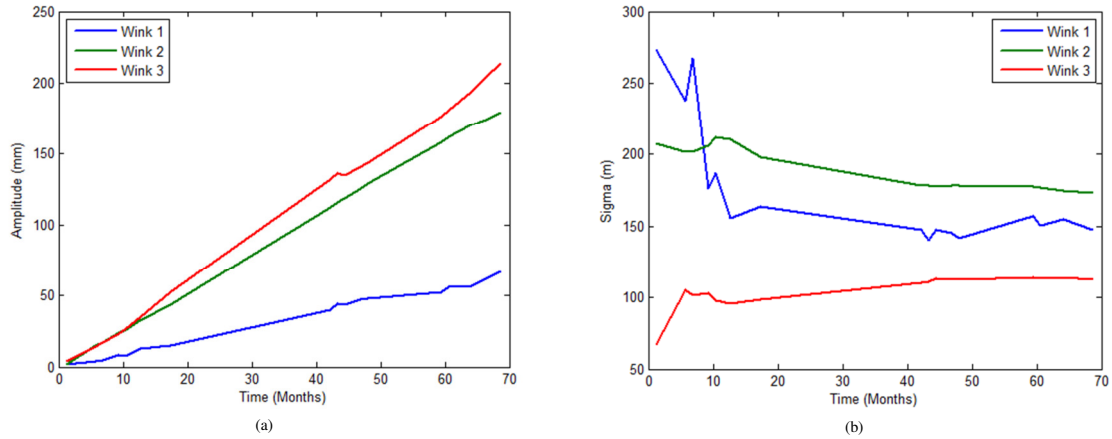


Figure 3-4: Temporal evolution of amplitude (a) and standard deviation sigma, σ_t , (b)

For all three cases, the amplitude displays a strong linear behavior (a), while the standard deviation parameter seems to jitter around for a bit then acquiesces to some asymptotic value (b). As mentioned previously, the Gaussian fits for about $t < 7$ months are somewhat suspect, as the data itself is still obfuscated by noise, and no obvious deformation trends have yet presented themselves.

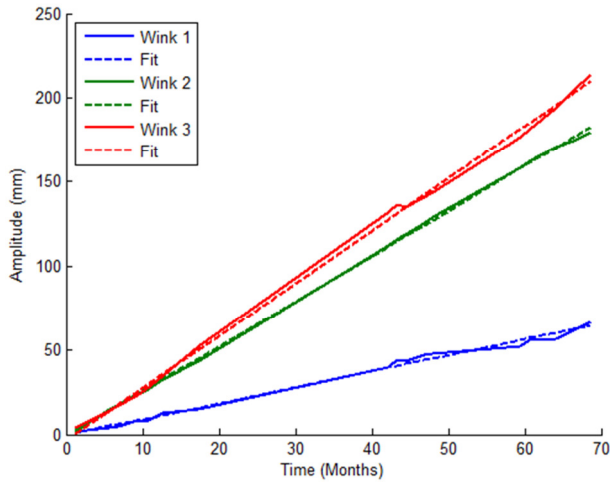


Figure 3-5: Temporal evolution of amplitude vs. linear regression

results are shown in Figure 3-5, and yield a correlation coefficient of nearly one for all three Wink Sinks.

Formalizing the observation of linearity in amplitude of the Gaussians over time,

$$\alpha_t = \alpha(t) = \alpha t \quad (3.2)$$

a normalized cross-correlation measure was once again employed upon the original amplitudes α_t and their

temporal linear regressions αt . The

To eliminate the standard deviation σ_t as a time-varying parameter, we repeated the above analyses with Gaussian fits generated with a fixed standard deviation σ —this standard deviation was chosen to provide the optimal fit across all times. Fortuitously, this yielded essentially identical results as compared to the time-varying standard deviation case.

Based on this preceding analysis, a spatiotemporal model for the displacement of a sinkhole was formulated: a Gaussian with a fixed standard deviation σ and linearly increasing amplitude. Thus the displacement incurred by a growing sinkhole centered at $\mathbf{x}_0 = (x_0, y_0)$ is given by

$$g(\mathbf{x}, t) = \alpha \cdot t \cdot \exp\left(\frac{-(\mathbf{x} - \mathbf{x}_0)^2}{2\sigma^2}\right) \quad (3.3)$$

This is the model used in the detection and classification algorithm which follows.

3.2 Method

The inherent spatial sparsity and non-uniformity of the data poses difficulties in the task of evaluating if, where, and to what extent subsets of the data agree with the postulated spatiotemporal sinkhole model. To address these issues, an algorithm was formulated which augments aspects of the Hough Transform with elements of matched filtering and residual measurement. We employ a parameter space search as in the Hough Transform; however, in lieu of the typically used ‘voting’ approach we instead make use of an approach similar to that in matched filtering—but the correlation operation is replaced with a residual-based measure. This blending allows for a fine control over parametric resolution, avoids the thresholding nature of Hough voting, and allows for one to robustly evaluate a match when presented with only a few data points.

Initially, we generate a parameter vector to index the desired parameter space. In the case of sinkhole detection, this amounts to accessing a four-dimensional parameter space with the vector $\mathbf{p} = [\mathbf{x}_0, \alpha, \sigma]$, where \mathbf{x}_0 is a 2d spatial parameter, α refers to a coefficient for the linear amplitude, and σ a fixed standard deviation. From here the algorithm proceeds as follows: for each point in the parameter space, \mathbf{p} , a corresponding spatiotemporal model $g_p(\mathbf{x}, t)$ is generated. The aptitude of this model is then evaluated with respect to the spatiotemporal point cloud $d(\mathbf{x}_i, t)$ over a region of influence $R(\mathbf{p})$, where i indexes a point contained within $R(\mathbf{p})$ and t refers to a time sample, via a residual $r(\mathbf{p})$ calculated from a metric μ .

In this application, a custom metric μ is used; this metric is based upon the notion of an absolute residual, i.e. the absolute difference between actual data values and their corresponding model predictions $res(\mathbf{x}_i, t) = |d(\mathbf{x}_i, t) - g_p(\mathbf{x}_i, t)|$. While one could simply use a standard metric such as the ℓ^1 or ℓ^2 norms (normalized by the number of points in (\mathbf{p})), these metrics would inherently be biased toward shallower (and less interesting) features. For example, given a sample dataset and a suitable fit $g_{\hat{\mathbf{p}}}(\mathbf{x}_i, t)$ the ℓ^1 norm might yield a residual \hat{R} . Simply scaling the dataset and fit by the same constant β will result in a residual $\beta\hat{R}$, even though the fit is intuitively just as suitable for the data as in the first case. Thus we desire a metric which yields amplitude scale-invariance in evaluating potential matches. To this end, we designed a metric μ which can factor in proportionality in quantifying a match:

$$\mu(\mathbf{x}_i, t) = \min\left(\frac{res(\mathbf{x}_i, t)}{\max(|d(\mathbf{x}_i, t)|, |g_p(\mathbf{x}_i, t)|)}, 1\right) \quad (3.4)$$

Here $res(\mathbf{x}_i, t)$ refers to the absolute residual, $d(\mathbf{x}_i, t)$ is the data point being compared, and $g_p(\mathbf{x}_i, t)$ is the value of the model at the same spatiotemporal coordinates (\mathbf{x}_i, t) . For the case where a data point and the model fit at that point have opposite signs, the measure is capped at a maximum value of 1. With this metric, we are able to accurately quantify the extent to which our displacement point cloud behaves like our sinkhole model in a spatiotemporal sense.

Because our model is by its definition Gaussian, we choose an influence region $R(\mathbf{p})$ which spans a circular region around the space being considered, \mathbf{x}_0 . The radius of this circle is designated by the fixed standard deviation in \mathbf{p} , and chosen to be 3σ ; this distance limits the scope of the residual calculation to relevant and significant data points, as the amplitude of the model beyond 3σ is less than 0.012% of the maximum.

Due to the non-uniform spatial distribution of the data, it falls into the realm of possibility that one might attempt to fit a spatiotemporal Gaussian specified by the model in a location where there is insufficient information for any measurements taken to be truly meaningful. For example, a ring of points which are present only on the outside of a Gaussian fit may lead to numerous spurious low-residual Gaussian fits, as the measure is only evaluated where there are data points in the first place. To address this informational disparity, we augment our calculation by partitioning the influence region $R(\mathbf{p})$ into subregions—each subregion has the measure μ applied separately and yields an independent residual.

$$R_1(\mathbf{p}) = \{\mathbf{x}_i \in d(\mathbf{x}_i, t) : 0 \leq \|\mathbf{x}_i - \mathbf{x}_0\| < \sigma\} \quad (3.5)$$

$$R_2(\mathbf{p}) = \{\mathbf{x}_i \in d(\mathbf{x}_i, t) : \sigma \leq \|\mathbf{x}_i - \mathbf{x}_0\| < 2\sigma\} \quad (3.6)$$

$$R_3(\mathbf{p}) = \{\mathbf{x}_i \in d(\mathbf{x}_i, t) : 2\sigma \leq \|\mathbf{x}_i - \mathbf{x}_0\| < 3\sigma\} \quad (3.7)$$

Thus in order to even use the measure in the first place, there must be data present in each significant region of the spatiotemporal Gaussian fit being considered. Once these partitions are created, we calculate an independent, normalized residual for each:

$$r_k(\mathbf{p}) = \frac{1}{(N_k \cdot T)} \sum_t \sum_{\mathbf{x}_i \in R_k(\mathbf{p})} \mu(\mathbf{x}_i, t) \quad (3.8)$$

N_k refers to the number of points in subregion $R_k(\mathbf{p})$, T is the number of time samples in the dataset, and t indexes into T . These k residuals are then combined via a simple average $r(\mathbf{p}) = [r_1(\mathbf{p}) + r_2(\mathbf{p}) + r_3(\mathbf{p})]/3$ yielding a four-dimensional matrix where each element is indicative of the extent to which the model $g_{\mathbf{p}}$ matches the data.

3.3 Results and Analysis

3.3.1 Validation—Simulated Data

Initially, the algorithm was tested upon simulated data in order to verify its functionality—both the case of sparse vs. uniformly-sampled data, as well as simply proving its ability to detect the model it sets out to find. To accommodate these goals, we generated three distinct datasets (Figure 3-6):

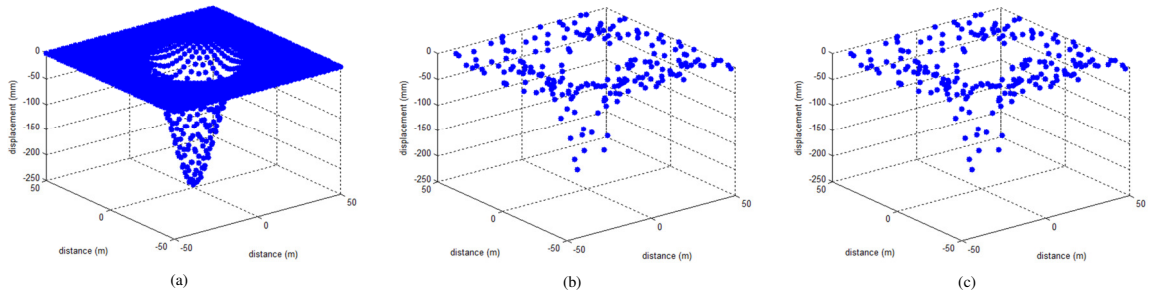


Figure 3-6: Simulated Datasets

- In the first set (a), we compute a spatiotemporal Gaussian which perfectly adheres to the previously determined sinkhole model (3.1). Here $\mathbf{p}_0 = [\mathbf{x}_0, \alpha, \sigma] = \left[(0\text{ m}, 0\text{ m}), -5.5 \frac{\text{mm}}{\text{month}}, 10\text{ m} \right]$. This instance of the model was sampled with a uniform grid spanning $\mathbf{x} = (x, y)$ where $-50\text{ m} \leq x \leq 50\text{ m}$ ($\Delta x = 2.5\text{ m}$) and $-50\text{ m} \leq y \leq 50\text{ m}$ ($\Delta y = 2.5\text{ m}$) which yielded 1681 spatial samples for each corresponding temporal sample.
- The second set (b), employs the same spatiotemporal Gaussian as described in (a); however, this model was instead sampled at 200 spatial locations selected from a bivariate uniform distribution $\mathbf{x} \sim \mathcal{U}_2(-50, 50)$
- The third set (c) utilizes a different model from (a) and (b), in that all temporal variation has been removed: $g(\mathbf{x}, t) = \alpha \cdot t_{10} \cdot \exp\left(\frac{-(\mathbf{x}-\mathbf{x}_0)^2}{2\sigma^2}\right)$. Succinctly, one time-sample has been isolated and duplicated for all times in T . In attempting to measure the spatiotemporal match of this dataset, it will still be matched against the original sinkhole model which varies with time (3.1).

Each of the three datasets was analyzed with our spatiotemporal matching algorithm, and a resulting four-dimensional residual $r(\mathbf{p})$ was calculated for every case. The parameter space $\mathbf{p}_0 = [\mathbf{x}_0, \alpha, \sigma]$ used in this analysis is as specified: $-50\text{ m} \leq x \leq 50\text{ m}$ ($\Delta x = 2.5\text{ m}$), $-50\text{ m} \leq y \leq 50\text{ m}$ ($\Delta y = 2.5\text{ m}$), $-10 \frac{\text{mm}}{\text{month}} \leq \alpha \leq \frac{10\text{ mm}}{\text{month}}$ ($\Delta \alpha = 0.25 \frac{\text{mm}}{\text{month}}$), and $2.5\text{ m} \leq \sigma \leq 30\text{ m}$ ($\Delta \sigma = 2.5\text{ m}$). In order to visualize and interpret the residual $r(\mathbf{p})$, we collapsed it from four dimensions into a two-dimensional form. Intuitively we are most interested in the ‘best case’ match at each spatial point in the residual \mathbf{x}_0 ; thus we take

$$m(x_0, y_0) = \min_{\alpha, \sigma} r(\mathbf{p}) \quad (3.9)$$

and term $m(x_0, y_0)$ the minimum residual image. The results are displayed in figure 3-7.

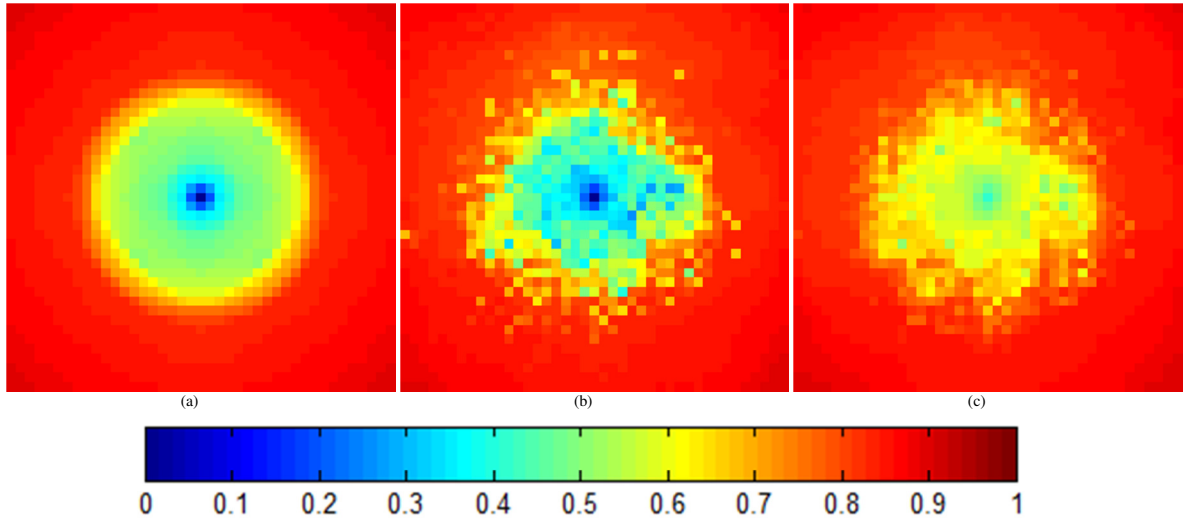


Figure 3-7: Minimum Residual Images

As anticipated, in case (a) we ascertain a perfect match when the model is perfectly overlaid with the template $m(x_0, y_0) = m(0m, 0m) = 0$. We also achieve a perfect match in the case of sparse, non-uniform sampling (b), with a minimum residual value of 0 in the same location. Interestingly, we still obtain somewhat of a response in the third case (c), with $m(0m, 0m) = 0.42$. This result is actually desirable, as it demonstrates the capacity for our matching algorithm to pick up on spatiotemporal data which adheres in a spatial sense, but not a temporal sense; additionally, it can handily differentiate between such data and that which adheres in a spatiotemporal way. Physically, in the context of detecting sinkholes, this is also a good thing—we achieve strongest results from those geohazards which evolve in a fashion most akin to our ground truth, but we also are able to detect (to a lesser extent) geohazards which move in more discrete motions.

For each case, the minimum-residual images also exhibit a response beyond the actual location of the ‘perfect’ overlay between the model and the data: in (a), a clear

radial fringing pattern presents itself. This is commensurate with lower-amplitude tangential fits between the data and smaller templates which are slightly offset from the center of the actual feature. This pattern also shows up in (b) and (c), but in less of a clear fringe. Some of the outlying pixels have somewhat stronger detections than those in (a), as the sparsity of the data allows for smaller (albeit still poorly-fitting) templates to slightly register in emptier areas.

In an effort to both further enhance visualization of the minimum residual and to accurately reflect the areal influence of a match, I here introduce the notion of the propagated minimum residual. Because the standard deviation parameter (σ_l) in our model inherently changes the footprint over which that model is applicable, a singular pixel in the original minimum residual $m(x_0, y_0)$ can correspond to a detection which encompasses and eclipses multiple other pixel residual values. Thus, for each spatial point in the minimum residual image we determine the corresponding σ which produced that value

$$\sigma_m = \operatorname{argmin}_{\sigma} \left[\min_{\alpha} r(\mathbf{p}) \right] \quad (3.10)$$

With this in hand, we then determine the corresponding influence region of each pixel in the minimum residual image—a circle of radius σ_m —overlap all of the regions, and take the minimum residual value at each pixel location from all of the overlays. The results of this operation are shown in figure 3-8.

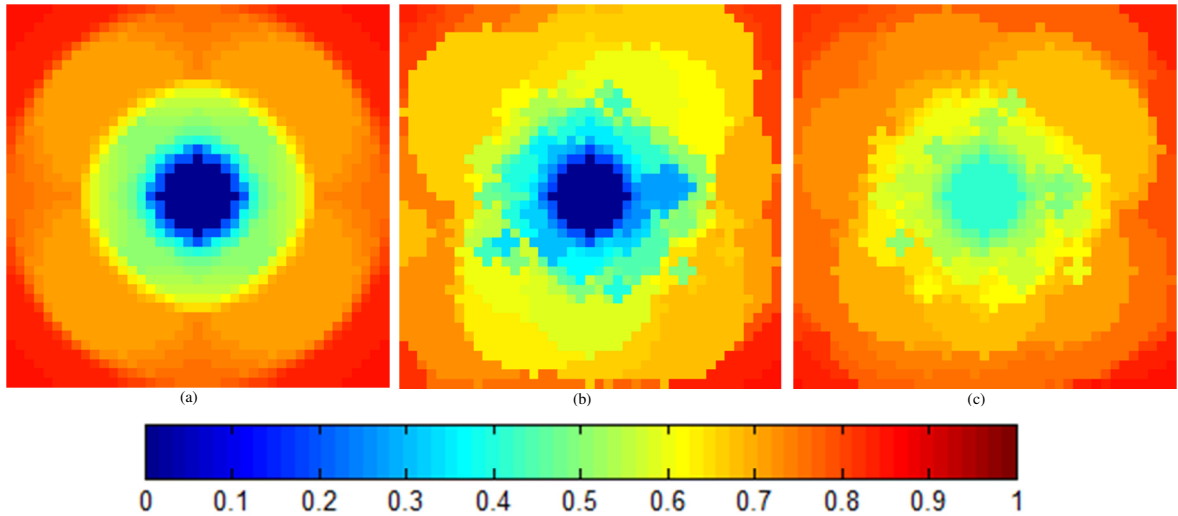


Figure 3-8: Propagated minimum residual images

For cases (a) and (b), this propagated residual image very clearly shows a strong detection which has roughly the same size as its corresponding displacement spatiotemporal Gaussian depression. Case (c) shows a similar region of influence, but for a diminished residual value.

These results on our simulated data verify the core functionality of our spatiotemporal matching algorithm, and illustrate that functionality remains effective in the case of sparse, non-uniformly distributed data. Additionally, this simulation also demonstrates the capacity our algorithm has for detecting behavior resembling our model, but with a diminished fit evaluation.

3.3.2 Wink Sinks

Next, the spatiotemporal matching algorithm was applied to the Wink Sink dataset, elaborated upon in section 2.3.1. Again, the spatiotemoral sinkhole model specified in (3.1) was employed as a matching template; the parameter space $\mathbf{p}_0 = [\mathbf{x}_0, \alpha, \sigma]$ here was chosen to encapsulate a broad range of potential sinkhole physical extents, as well as a healthy span of sinkhole deformation rates of interest. Thus we selected: $0m \leq x \leq$

$6,670m$ ($\Delta x = 10m$), $0m \leq y \leq 8290m$ ($\Delta y = 10m$),

$-5 \frac{mm}{month} \leq \alpha \leq -1 \frac{mm}{month}$ ($\Delta \alpha = 0.25 \frac{mm}{month}$), and $5m \leq \sigma \leq 185m$ ($\Delta \sigma = 10m$). The

resulting propagated minimum residual is shown in Figure 3-9a. While this image is informative as is, in analyzing an actual dataset it is desirable to make a distinction between low-residual detections which subside at a large rate, and those which subside more slowly. Thus we employ a straightforward mapping upon the residual r

$$\rho(\mathbf{x}_0, \alpha, \sigma) = [1 - r(\mathbf{x}_0, \alpha, \sigma)] \exp\left(\frac{1}{\alpha}\right) \quad (3.11)$$

In this way, we prioritize first low-residual matches, then further those low-residual matches which are subsiding at a faster rate. By virtue of this mapping, high-risk matches now correspond to maxima in the risk matrix ρ . This risk mapping can be visualized in a manner analogous to the residual visualization, i.e. collapsing ρ into its spatial dimensions (as in Equation 3.9, but with a maximum operation instead) to generate a maximum risk image, and using a propagation technique (via Equation 3.10, but also with maximum operations) to accurately portray the areal influence of risk detections. Figure 3-9b shows the resulting propagated maximum risk image for the Wink dataset; here one can see that some regions which displayed a low-residual in Figure 3-9a have attenuated risk values due to their low subsidence rates. However, more significant detections retain a large risk value.

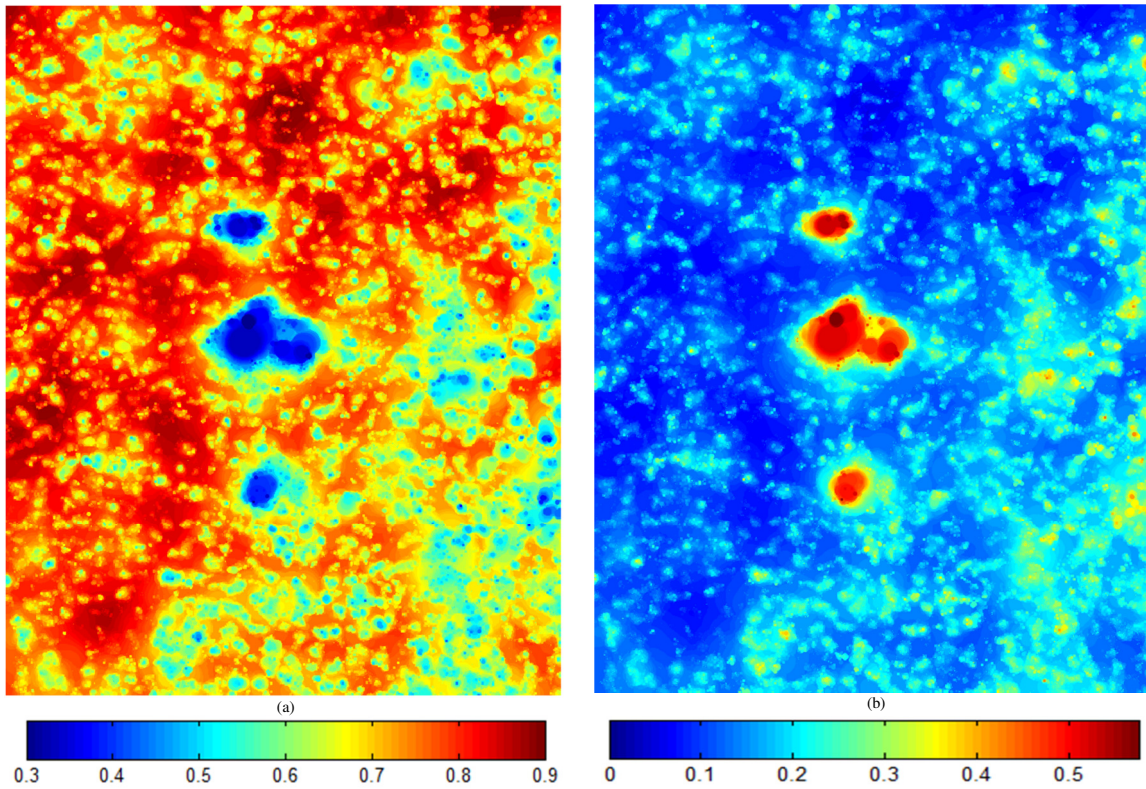


Figure 3-9: Propagated minimum residual image for the Wink dataset (a), corresponding propagated maximum risk image (b)

Applying a binning similar to that used in the display of graph cut detections (2.3.1) upon the propagated maximum risk image yields detections shown in Figure 3-10. As before, the binning encompasses three visible detection rankings—red is indicative of severe risk, orange with moderate risk, and yellow with slight risk. As can be seen, our spatiotemporal matching algorithm handily locates all four of the prominent Wink geohazards; additionally, several smaller detections present themselves as well



Figure 3-10: Binned propagated maximum risk overlaid upon visual satellite imagery

(predominantly on the right side of the image). Specifically, 100% of Wink 1 is detected—52.6% is labeled as severe and 47.4% is labeled as moderate; 100% of Wink 2 is detected—81.6% is labeled as severe and 18.4% is labeled as moderate.

4. Approach Comparison

Both of the methods introduced, the graph theoretic method and the parametric spatiotemporal matching method, display efficacy in locating sinkholes. Because they are entirely independent processes, it raises the question whether one might supplant the other in use, or if they can exhibit complementary functions. This section compares various aspects of each method and determines an answer to the previously posed question concerning method necessity and exclusivity.

4.1 Simulated Data

In an initial foray into determining superiority, I revisit the case of simulated data, first presented in section 3.3.1. As elaborated before, this data consists of three distinct sets (Figure 4-1): a grid-sampled, spatiotemporal Gaussian with amplitude that evolves linearly in time; a randomly-sampled version of the exact same model; and a randomly sampled spatiotemporal Gaussian which exhibits a constant amplitude with respect to time.

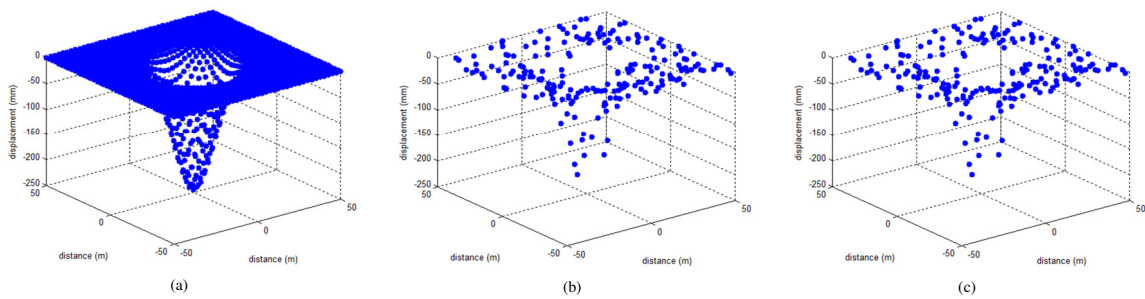


Figure 4-1: Simulated datasets

First, I apply the graph cut sinkhole detection algorithm upon each of the three datasets. In Figure 4.2, a-c show the nodal scores which result. For case I, the graph cut algorithm displays very low (ergo good) scores (a), and the resulting segmentation (d)

very cleanly outlines the spot as expected. The graph cut algorithm is able to achieve similar results in the case of sparse data (b); this is not surprising, as the method is non-Cartesian at its core and only operates upon relational neighborhood characteristics. However, because the algorithm attempts to translate nodal detections into Cartesian regions as its final output, the resulting segmentation in (e) is not quite as clean as that achieved in the case of dense, grid-based data. In the third case, i.e. a spot which does not evolve with time, the graph cut algorithm is extremely harsh. This makes sense, as 75% of the calculated score is obtained from time-windowed linear regressions of displacement—in this case, those linear regressions yield lines with a zero-slope for the majority of the temporal windows. Thus, there is no displacement change upon which to operate, and nearly the maximum (worst) score is assigned for this temporal component. As a result, there is no sinkhole segmentation (f) whatsoever.

Next I evaluate the dataset for second sinkhole detection algorithm, spatiotemporal matching. In order to avoid repeating results from 3.3.1, I here calculate the propagated maximum risk extracted from each dataset (employing the spatiotemporal Gaussian in Equation 3.3, established as our sinkhole model). These propagated risk images are shown in Figure 4.2, g-i. Reassuringly, they are consistent with the conclusions garnered earlier from the analysis of the propagated minimum residual images. Additionally, I bin these propagated risk images and display the resulting risk segmentations (i-l). It is clear that the spatiotemporal matching detector is able to operate with more fidelity in sparse conditions, as its detections more accurately reflect the areal influence and implications of the data. Also, the spatiotemporal algorithm displays a lone ability to detect sinkholes

which evolve in a more discrete fashion time-wise, and with a lesser score as well. The graph cut algorithm fails completely in this case.

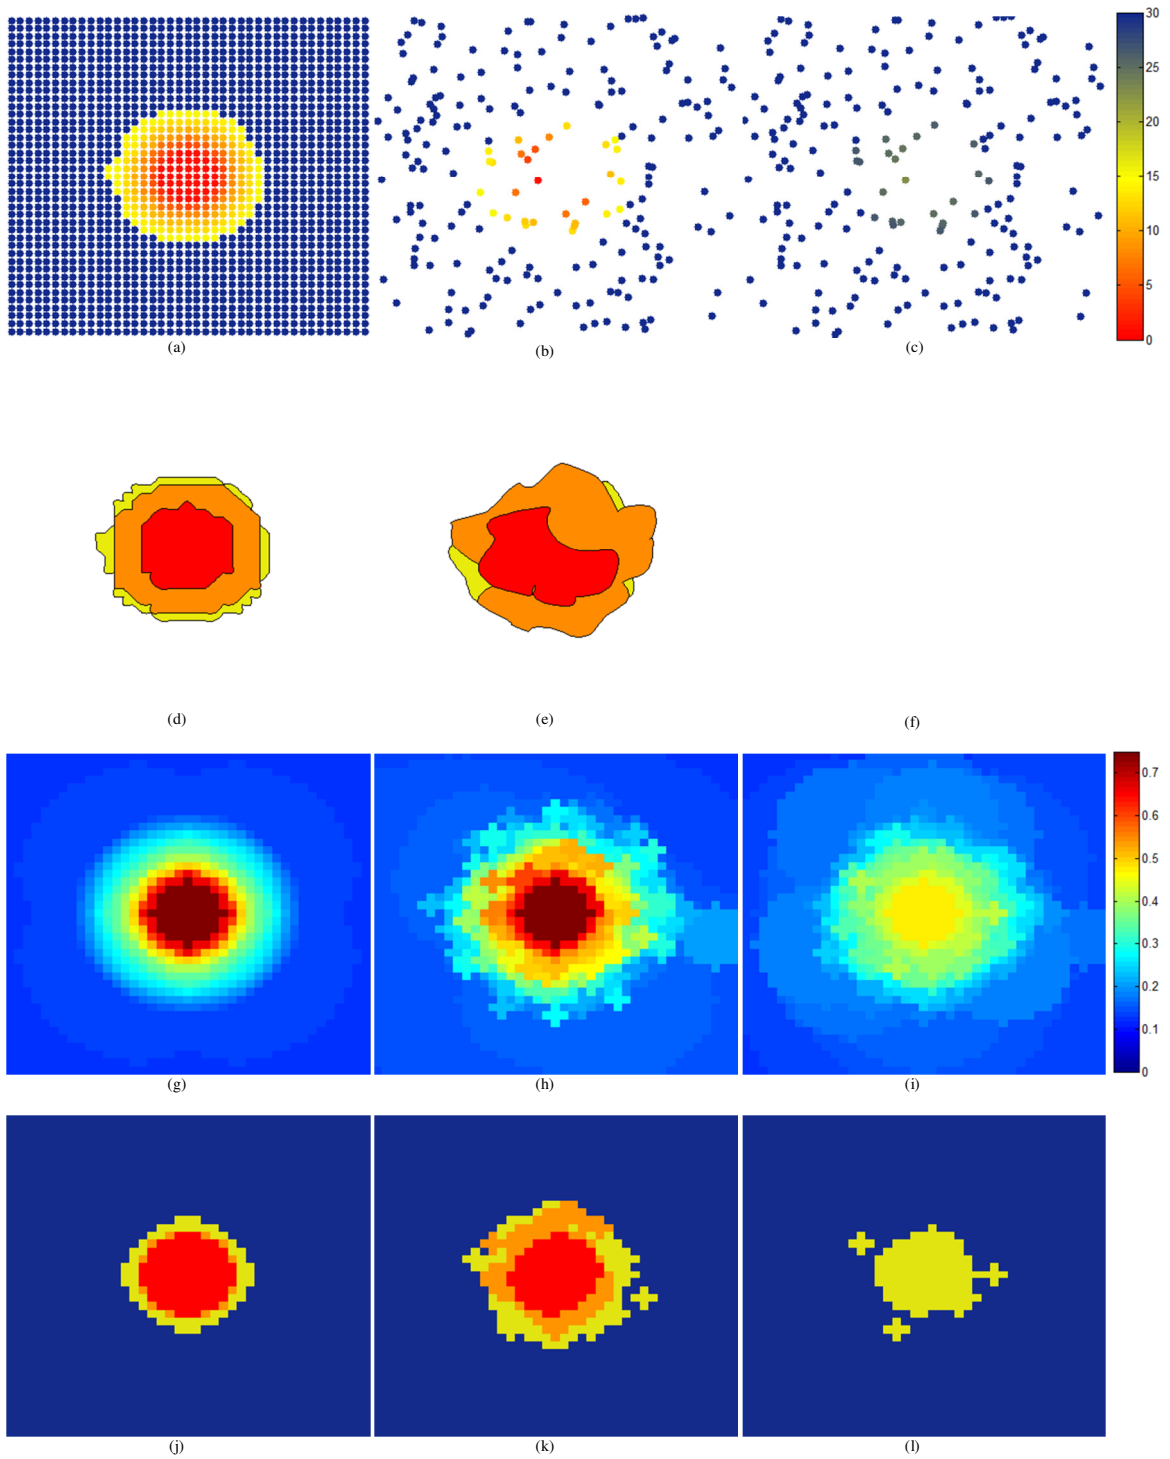


Figure 4-2: Comparison of the two methods on simulated dataset

4.2 Wink Dataset

In addition to revisiting the simulated dataset, I here re-employ the Wink dataset with which both algorithms have seen use. This is the only dataset for which we have ground truth—albeit limited—and thus provides a foundational metric upon which comparisons can be made between the two methods. The general segmentations are shown side-by-side in Figure 4-3, where (a) displays the segmentation obtained from the graph cut routine, while (b) gives that generated by the spatiotemporal matching algorithm. As can be seen, both approaches identify the same general areas as threats—particularly, the four Wink geohazards. However, the graph cut method (a) looks to have a slightly higher sensitivity for peripheral detections (large, yellow areas on the right side of the image), but at the expense of accurately portraying the severity of the main four detections. This downfall is handily covered by the matching algorithm.



Figure 4-3: Resulting segmentation of the Wink dataset for the graph cut algorithm (a) and the spatiotemporal matching approach (b).

This aforementioned coverage is well-illustrated in Figure 4-4, which immediately follows. This figure shows the two collapsed sinkholes, Winks 1 and 2, with their respective segmentations from both algorithms. Neither segmentation of Wink 1 is particularly satisfactory—this is attributed to the lack of data in the immediate vicinity of the sinkhole, which collapsed before the data was acquired. As elaborated previously, the graph cut algorithm labels 37.5% of Wink 1 as severe, 57.8% as moderate, and 4.7% as slight. The matching approach has a superior detection, labeling 52.6% of Wink 1 as severe and 47.4% as moderate. Wink 2, however, clearly sets the two algorithms apart in terms of detection coverage and accuracy. One might argue this is the most important of the two ground truth sinkholes, as it collapsed shortly after the dataset was acquired, and displayed the most pertinent and immediate characteristics of a collapsing sinkhole. The graph cut method manages to squeeze only 17.2% of the sinkhole into the severe category, with the remaining 82.8% pinned as moderate. The matching technique excels, attributing 81.6% of Wink 2 as a severe detection and a mere 18.4% as moderate.

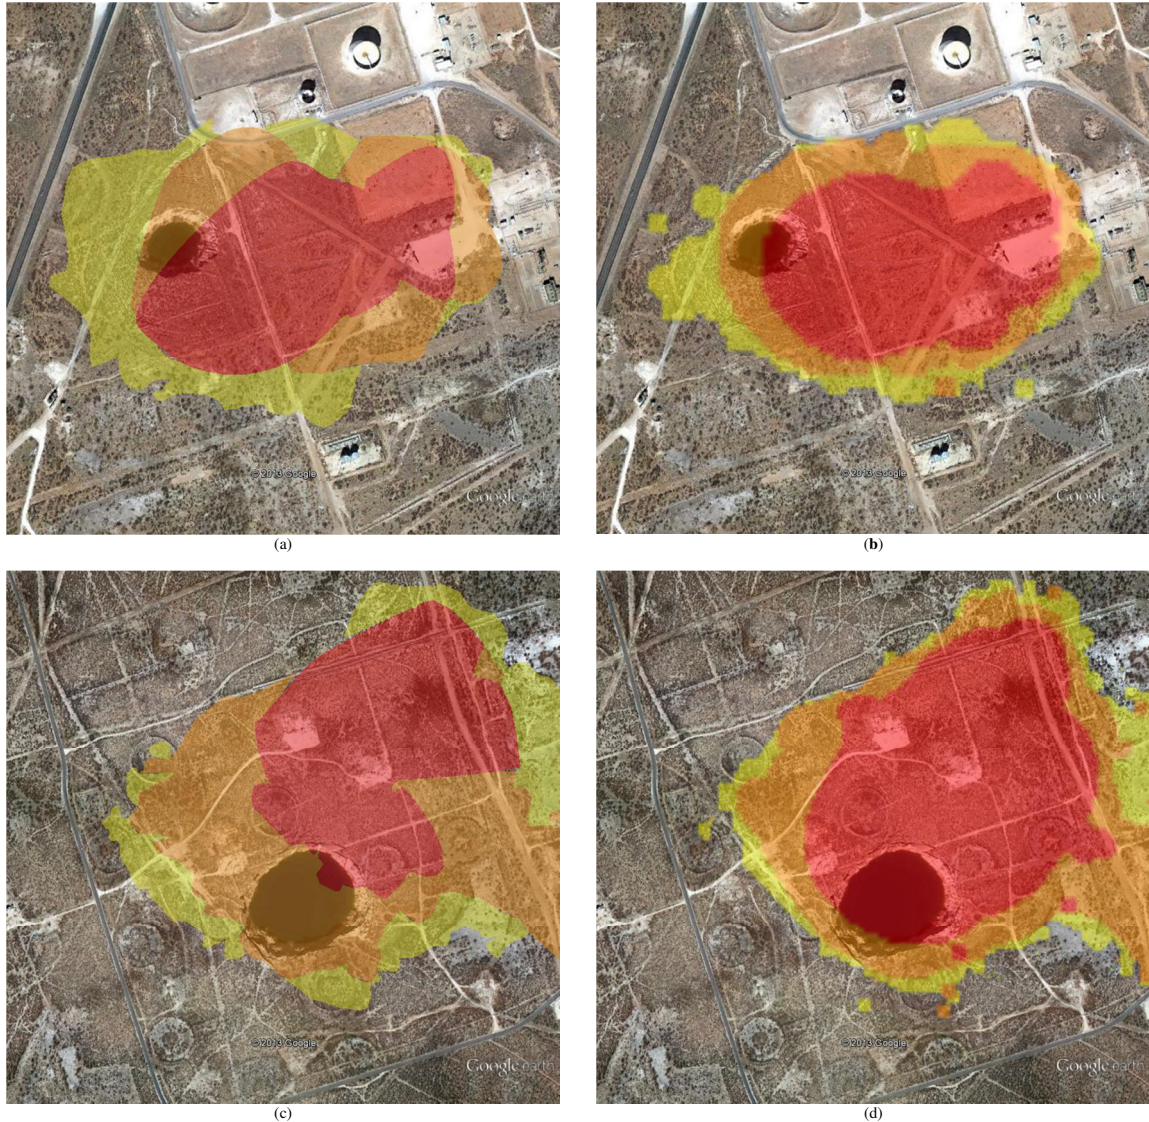


Figure 4-4: Detailed view of Winks 1 and 2 (top and bottom rows) for graph cut algorithm (a,c) and spatiotemporal matching (b,d)

4.3 Virginia Data—General Comparisons

As part of our study, we also ordered the acquisition of InSAR images for a small portion of Virginia. Specifically, this area spanned 160 square kilometers around the city of Staunton, and 32 SAR scenes were obtained from August 29, 2011 through October 25, 2012. This InSAR stack was processed with the PSInSAR/SqueeSAR processing suites, and yielded 296,121 spatial data points with 31 temporal displacement samples.

Both algorithms were used to analyze this Virginia data, resulting in the detections shown (Table 4-1). For the latter algorithm, the parameter space was chosen to span the data spatially, and provide coverage of various sinkhole sizes and deformation rates:

$$0m \leq x \leq 43,235m (\Delta x = 2.5m), 0m \leq y \leq 51,768m (\Delta y = 2.5m), -5 \frac{mm}{month} \leq$$

$$\alpha \leq -1 \frac{mm}{month} \left(\Delta \alpha = 0.25 \frac{mm}{month} \right), \text{ and } 5m \leq \sigma \leq 100m (\Delta \sigma = 5m). \text{ The base images}$$

yield a spatial resolution on the order of a couple of meters per pixel.

Algorithm	Total Detections	Severe Detections	Moderate Detections	Slight Detections
Graph Cut	675	0	503	170
Matching	793	7	124	598

Table 4-1: Overall detection distributions

In terms of sensitivity, both algorithms seem to result in a similar number of overall detections, with the matching algorithm yielding slightly more segmented regions. However, the severity distribution of detections is vastly different between the two. A whopping 74.7% of the detections yielded by the graph cut approach are classified as moderate, while the remainder are classified as slight risk. This provides a stark contrast to results from the matching algorithm, which is more heavily predisposed to slight detections (82.0% of results) but does offer higher-confidence, severe detections. This departure between the two algorithms is surprising, as the risk binning for both are calibrated using the Wink dataset; thus one would expect more consistent results from both. Potentially, this is to be expected from the fundamental differences between the two algorithms. The graph cut algorithm performs more admirably in particularly data-sparse areas—it requires that there be neighboring data points a reasonable distance away which support a subsiding, spot-like detection, while the matching algorithm is more stringent in its requirements for the presence of data. This most likely explains why there are

numerous, small moderate sinkhole detections from the first approach. The contrast in sensitivity to slight detections is most likely due to the aptitude of the matching approach in evaluating partial matches to the prescribed sinkhole model. As shown and discussed in Figure 4-2, the graph cut approach does not display the same level of nuance in quantifying borderline detections.

Lastly, I discuss concurrency between the detections of the two methods—does one or the other provide coverage for detections made by the alternative, or do both yield unique information? To this end, the spatial commonalities were measured with respect to the segmentations from both approaches—overlap for a given segmentation is approximated for the case where a segmentation for a corresponding algorithm occluded by at least 10% of the area.

Algorithm	Total Overlap	Severe Risk Overlap	Moderate Risk Overlap	Slight Risk Overlap
Graph Cut	29.8%	N/A	31.4%	25.3%
Matching	20.7%	57.1%	52%	18.1%

Table 4-2: Detection overlap between methods

From these results (Table 4-2), it looks as if both methods provide valuable information; in particular, the matching algorithm provides more uniqueness (in addition to far more detections) for borderline segmentations, and the graph cut routine yields less redundant moderate detections. Additionally, many of the severe risk detections achieved from the matching approach are corroborated by lesser-risk conclusions from the graph cut implementation; the same is true for moderate segmentations made by the matching routine.

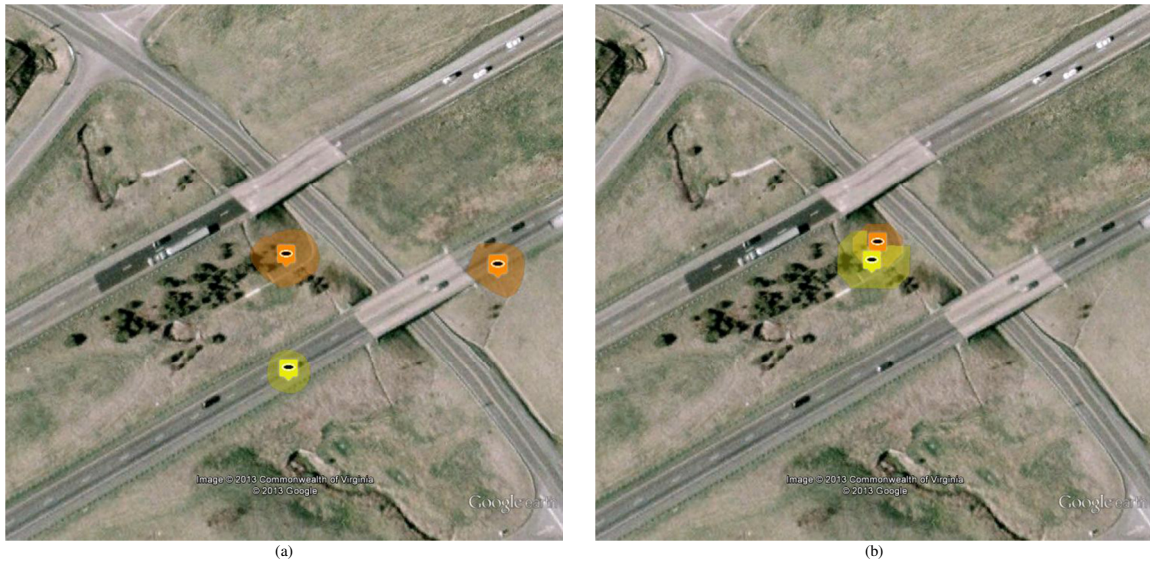


Figure 4-5: Example overlap and uniqueness between graph cut detector (a) and matching algorithm (b)

As illustrated in Figure 4-5, these detections can often line up quite well, substantiating results from the corresponding algorithm. Additionally, the absence of a detection from a competing algorithm can result in the loss of potentially worthwhile information. Thus, this analysis of concurrency between the two methods leads to two distinct conclusions: one can veer towards a conservative standpoint by taking the union of detections from both algorithms—this would increase the number of false positives, but also lower the number of missed detections. Alternatively, one could minimize the need for extensive in-field surveying and verification by taking the intersection of these detections—this would have the opposite effect, lowering false positives but increasing the number of missed sinkholes.

5. Conclusions

Both of the methods deliberated upon in this thesis present viable options in the pursuit of automatic sinkhole detection. Moreover, as discussed in the previous section they offer complementary functionality, with the spatiotemporal matching approach demonstrating a marked superiority in its detections.

Qualitatively, the main qualm leading to this gap in efficacy lies in the graph cut detector's lack of a hard model upon which to base detections. This detector is truly a combination of sparse spot detections, and because it operates on a relational graph structure these spots can be quite amorphous in shape. Essentially, the process amounts to locating regions which are subsiding in a spot-like fashion, and gives more credence to those regions which are subsiding more significantly and in a more coherent fashion. However, this shortfall is completely made up for with the second sinkhole detection method, spatiotemporal matching. The foundation of the algorithm is based entirely upon validation of a prescribed model—adherence to this model in a spatiotemporal sense results in a strong detection. There is one caveat to be mentioned here; this method is only as potent as the model in question is accurate.

Given the results from Section 4.3, I contend that neither method be discarded in the endeavor to detect sinkholes; instead, I think the most prudent route is to combine the segmentations from both algorithms. Thus we can exploit the various advantages of each method—the graph cut algorithm deals better with very sparse conditions, the matching algorithm can pick up on nuanced detections, etc.—but at the same time give priority to the demonstrated superior approach where overlaps do occur. Additionally, this allows

for corroboration between methods, which could perhaps raise the risk factor for areas which are verified as such.

5.1 Applicability/Generality

If one seeks to apply a similar approach to a different problem which can be measured with ground deformation, the work in this thesis—specifically, the parametric spatiotemporal matching approach—can be adapted to suit the needs of this new problem. Namely, this is accomplished by virtue of the generality of the method such that any potential model might be employed (with an implementation alteration to the region partitioning used for the sinkhole model) in determining a spatiotemporal match. The same parametric search and residual generation can be leveraged, potentially with a custom-defined risk mapping as well.

Additionally, the model generality provided by the algorithmic framework provides a measure of adaptability for our original problem, sinkhole detection. Should our originally prescribed model prove inadequate in any way, the matching algorithm offers flexibility in refining/replacing this model via analysis from field verification, or perhaps a computational simulation involving soil/rock mechanics and interactions with water activity.

Bibliography

- [1] W. Zhou and B.F. Beck, "Management and mitigation of sinkholes on karst lands: an overview of practical applications," *Environmental Geology*, vol. 55, pp. 837-851, 2008.
- [2] F. Gutierrez, A.H. Cooper, and K.S. Johnson, "Identification, prediction, and mitigation of sinkhole hazards in evaporate karst areas," *Environmental Geology*, vol. 53, pp. 1007-1022, 2008.
- [3] A. K. Gabriel, R. M. Goldstein, and H. A. Zebker, "Mapping small elevation changes over large areas: differential radar interferometry," *Journal of Geophysical Research*, vol. 94, no B7, pp. 9183-9191, July 1989.
- [4] A. Ferretti, C. Prati, and F. Rocca, "Permanent scatterers in SAR interferometry," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 1, pp. 8-20, Jan. 2001.
- [5] A. Ferretti, A. Fumagilli, F. Novali, C. Prati, F. Rocca, and A. Rucci, "A new algorithm for processing interferometric data-stacks: SqueeSAR," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 9, pp. 3460-3470, Sept. 2011.
- [6] J.G. Paine, S.M. Buckley, E.W. Collins, and C.R. Wilson, "Assessing collapse risk in evaporate sinkhole-prone areas using microgravimetry and radar interferometry," *Journal of Environmental and Engineering Geophysics*, vol. 17, No. 2, pp 75-87, June 2012.
- [7] M.L. Rucker, B.B. Panda, R.A. Meyers, and J.C. Lommler, "Using InSAR to detect subsidence at brine wells, sinkhole sites, and mines," *Carbonates and Evaporites*, vol. 27, pp. 1-7, Feb. 2013.
- [8] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888-905, Aug. 2000.
- [9] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124-1137, Sep. 2004.

- [10] X. Lin, B. Cowan, and A. Young, "Model-based graph cut method for segmentation of the left ventricle," *IEEE Eng. In Medicine and Biology*, vol. 27, pp. 3059-3062, Jan. 2006.