

Received October 12, 2019, accepted October 31, 2019, date of publication November 5, 2019, date of current version November 15, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2951605

Subsurface Boundary Geometry Modeling: Applying Computational Physics, Computer Vision, and Signal Processing Techniques to Geoscience

RAYMOND LEUNG , (Member, IEEE)

Australian Centre for Field Robotics, The University of Sydney, Sydney, NSW 2006, Australia

e-mail: raymond.leung@sydney.edu.au

This work was supported by the Australian Centre for Field Robotics and the Rio Tinto Centre for Mine Automation.

ABSTRACT This paper describes an interdisciplinary approach to geometry modeling of geospatial boundaries. The objective is to extract surfaces from irregular spatial patterns using differential geometry and obtain coherent directional predictions along the boundary of extracted surfaces to enable more targeted sampling and exploration. Specific difficulties of the data include sparsity, incompleteness, causality and resolution disparity. Surface slopes are estimated using only sparse samples from cross-sections within a geological domain with no other information at intermediate depths. From boundary detection to subsurface reconstruction, processes are automated in between. The key problems to be solved are boundary extraction, region correspondence and propagation of the boundaries via contour morphing. Established techniques from computational physics, computer vision and signal processing are used with appropriate modifications to address challenges in each area. To facilitate boundary extraction, an edge map synthesis procedure is presented. This works with connected component analysis, anisotropic diffusion and active contours to convert unordered points into regularized boundaries. For region correspondence, component relationships are handled via graphical decomposition. FFT-based spatial alignment strategies are used in region merging and splitting scenarios. Shape changes between aligned regions are described by contour metamorphosis. Specifically, local spatial deformation is modeled by PDE and computed using level-set methods. Directional predictions are obtained using particle trajectories by following the evolving boundary. However, when a branching point is encountered, particles may lose track of the wavefront. To overcome this, a curvelet backtracking algorithm has been proposed to recover information for boundary segments without particle coverage to minimize shape distortion.


INDEX TERMS Active contours, backtracking, contour morphing, directional prediction, particle trajectories, spatial correspondence, subsurface boundaries, wavefront propagation.

I. INTRODUCTION

This paper considers the feasibility of modeling geospatial boundaries using differential geometry given sparse observations. The objective is to extract surfaces from spatial patterns and obtain coherent directional predictions along the boundary of the extracted surfaces. Modeling underground geological formations is challenging in general because the measurements are sparse and indirect. Due to operational constraints and the significant costs associated with data

gathering,¹ the available observations may not paint a complete picture in terms of spatial coverage. These measurements, although point-based, differ from those encountered in computer vision or image processing in some significant ways. The input consists of sparse spatial patterns in the form of irregularly spaced labeled drilled holes. Not only are the measurement locations sparse in the x-y plane, the sampling

¹This includes operator cost, energy expenditure for drilling, replacement cost of mechanical parts, efficiency cost of coordinating dependent processes such as blasting and excavation, and the cost of performing chemical assays to determine the composition, material type or geological domain associated with each sample.

The associate editor coordinating the review of this manuscript and approving it for publication was Yizhang Jiang .




FIGURE 1. Processing pipeline – a high level abstraction of the Boundary Geometry Modeling (BGM) framework.

is less dense along the z -axis. Geo-assay data is typically collected sequentially on a bench-by-bench basis from the top down (each bench has a height of $\sim 10\text{m}$) with significant time lapse in-between. In contrast, pixels are captured almost instantaneously using an image sensor array. The combination of these two factors means volumetric segmentation approaches that utilize 3D partial derivatives, and those that conduct minimal path search with respect to z , are not applicable as there are no voxels or fine-grain data available at intermediate depths between successive benches.

From a system perspective, the output provides a volumetric reconstruction of subterranean surfaces that conveys directional information. The motivation for predicting the slope of extracted boundaries is to provide guidance for more targeted drilling and exploration. The accuracy of this directional information needs only be commensurate with the resolution of the raw input (roughly $\sim 5\text{m}$) for it to be useful in a mining context. However, the directional estimates need to be coherent. An example of what not to do is using the outward normals of a contour as a means for extrapolation which inevitably cross-over when non-convex boundaries are involved. Hence, partial differential equations (PDE) are used to describe boundary movement in a more principled manner.

The input data used in this work contrasts with dense data sources such as point cloud produced by terrain laser scanners (LIDAR) [1], [2] and high-resolution slices generated by computed tomography [3] and presents its own challenges. Data incompleteness, sparsity, causality and resolution disparity are some of the issues to contend with and probable reasons for why differential geometry has not been more widely used in subterranean geology modeling. Most established PDE techniques for modeling surfaces operate on uniform grid data whereas our input samples are irregularly spaced. To overcome this, a bridging step based on boundary detection and edge map synthesis is described. This enables level-set methods to be applied to contours extracted from sparse non-uniform data points.

As an overview, Fig. 1 shows the processes for achieving the ultimate goal. Once a set of contours emerge following boundary extraction, the work flow next enters the spatial correspondence (reasoning) phase whose purpose is to associate ‘source’ regions with ‘target’ regions at successive intervals and estimate component displacements. This problem shares many similarities with object tracking in computer vision, but is made difficult by significant variations rather than gradual changes in contour shape. Often, the vertical resolution is low,

whilst some regions may not be sampled adequately, or at all, due to operational constraints. A region association and translation estimation approach is proposed to deal with the complexity of region merging and splitting from a resource allocation perspective.

Each set of associated source–target contours define a region in a motion-compensated frame for which contour morphing (spatial warping) [3] is applied. The objective is to model the shape of the boundary in between two cross-sections as a propagating interface. The underlying premise is that local surface deformation can be described as an evolutionary process governed by some PDE. Although the exact form used in this work might not match reality in terms of ore genesis, it is a reasonable alternative to unconstrained warping approaches and does provide a continuous deformable model. Using level-set methods [4], topological changes can be handled seamlessly during the morphing process. To facilitate slope prediction, the evolving interface (contour boundaries) are tracked using particle trajectories. This works well along portions of the boundaries where differences in curvature between the source and target contours are small. It breaks down when branching occurs, i.e., when a curvelet emanates from a single point. To address this, a novel backtracking algorithm has been proposed to recover lost information during particle advection and minimize boundary distortion attributed to tracking failure.

A. RELATED WORKS

This paper is distinguished from prior work through its attempt in fully automating a chain of processes required to reconstruct subterranean surfaces using sparse labeled data. These processes include boundary extraction, spatial correspondence and contour evolution, as outlined in Fig. 1. Although the contours for various geological domains (henceforth referred as *geozones*) can be specified interactively, our model basically requires only a set of unordered points, sampled non-uniformly across an orebody in multiple cross-sections as input.

In [5], Sprague and de Kemp presented a partially automated tool which uses Bézier and NURBS (non-uniform rational B-spline) curves to model non-planar 3D surfaces. Fitting under the guidance of a control frame, it makes use of 2D *interpreted* plan views provided by mine geologists, created using local slices of projected drill hole data at semi-regular spaced depths. The authors emphasized that in moving from 2D polygonal lines to 3D surface construction,

the continuation of common features requires delicate spatial synchronization across neighboring plan-sections, and manual correspondence performed by the geologist was critical to the integrity of the modeled structure and preventing self-intersections. The term *map trace* was defined as a “geological interpretation delineating the intersection of a geologic boundary as it breaks through the surface... or as it intersects a given elevation plane”. This definition closely resembles our notion of *active contour extracted boundaries* which represent segmented regions in a geozone. Their techniques were refined by imposing positional and orientation constraints using structural ribbons (expert knowledge), thus it may be categorized as a semi-supervised technique.

In [6], Mitášová and Hofierka applied differential geometry, more specifically, regularized spline with tension to topographic analysis of a watershed. This utilized convex/concave sections with smooth curvature for interpolation. However, the focus was to model the top surface (as opposed to sub-surfaces) using dense elevation data obtained via remote sensing (as opposed to sparse data).

In [7], Kaufmann and Martin built 3D subsurface models using a variety of sources (drill holes, cross-sections and geological maps) with different motivations. Their goal was to further understand subsoil characteristics such as hydro-geologic or geothermic properties of the geological bodies. Their surfaces were modeled using DSI (discrete smooth interpolation) which computes the location of nodes by balancing roughness and misfit constraints (see Mallet [8] and Frank [2]). This is representative of a class of information-rich GIS fusion approaches which utilize topographic, geological and structural data [9]. In contrast, our approach makes the most of the limited data in an information-poor environment where only blasthole locations and geozone labels are available.

In [10], Caumon et al. presented general guidelines for creating a 3D structural model made of faults and horizons using sparse field data. Their focus was natural resource evaluation and hazard assessment; triangulated surfaces with variable resolution was also discussed. The authors offered many insights, one notable comment is that “3D subsurface modeling is generally not an end, but a means of improving data interpretation through visualization... to generate support for numerical simulations of complex phenomena (i.e., earthquakes, fluid transport) in which structures [11] play an important role.” This is also true of subsurface models serving as decision support tools in mining and geological exploration.

In [12], Dirstein et al. demonstrated that automated surface extraction and segmentation of peaks and troughs from seismic survey can provide insight into structural and sedimentary morphology. In particular, differential geometry was used to select objects of concave and convex curvature, these features can help identify subtle cues for fluid flow events that are perhaps over-looked by conventional interpretation methods. For an in-depth survey of past efforts and current interests in 3D geological mapping, readers are referred to

[13]–[15] and [16]. Relevant techniques from computational physics and computer vision are presented in Appendixes A and B. As a quick overview, the major themes and related works are highlighted in Table 1.

TABLE 1. Computational physics and computer vision techniques used in this work (refer to Appendixes A–B in the Supplemental Material).

	Description and referenced works
A	Theoretical treatment of active contours & gradient vector field (Kass et al., 1988) [17], (Caselles et al., 1997) [18], (Ivins and Porrill, 1995) [19], (Xu and Prince, 1998) [20], (Horn and Schunck, 1981) [21].
B	Foundations for contour metamorphosis (Nilsson et al., 2005) [3], (Breen and Whitaker, 2001) [22], (Museth et al., 2005) [23], (Nielsen and Museth, 2006) [24], (Bertalmio et al., 2000) [25], (Osher and Sethian, 1988) [26], (Peng et al., 1999) [27], (Nagashima et al., 2007) [28]

In both Dirstein’s and our work, curvature preserving spatial representations are used to good effect albeit the objectives and modalities are different. One aspect of Dirstein’s work is waveform analysis, where genetic fitness (similarity) relative to a genotype (waveform signature) can reveal relative stability of neighboring segments above and below a particular surface and confer understanding about its structure and stratigraphy. In our work, we obtain an essentially continuous representation for the slope of geozone subsurfaces from sparse boundary patterns through contour extraction and metamorphosis; this provides useful directions for subsequent drilling and exploration. Our motivation stems from an interest in integrating techniques from pattern analysis, computational physics, computer vision and signal processing, bringing these to bear on a subsurface reconstruction problem, solving it in unconventional ways.

B. APPLICATION SCENARIO

The target application is the modeling of subsurface boundaries in an open-pit mine of sedimentary iron ore deposits.² Here, the sparse spatial patterns derive from blasthole samples where each is labeled with a geozone. The *geozone* labels provide a classification of geological domains based on mineralization and stratigraphic units [29]. From a mining perspective, these are used to segregate high-grade minerals from low grade or unprofitable materials. With some effort, the geozones can be turned into coherent spatial clusters. However, as unorganized points, or even with a discretized block-based representation, there is no easy way of obtaining a coherent motion field, one that describes the direction a boundary moves in consistently without artifacts or discontinuities. This is one major reason for using a curve-based, deformable surface model.

At the outset, it is important to distinguish this work from 3D segmentation approaches [30], [31] that use energy

²The most common minerals are hematite $\text{Fe}_2^{3+}\text{O}_3$ and goethite $\text{Fe}^{3+}\text{O}(\text{OH})$. These deposits are believed to have formed as chemical precipitates on the floor of shallow marine basins in a highly oxidizing environment during the Proterozoic eon, circa 1.9–2.4 billion years ago.

functional for minimum path optimization given contours at specific depths which is a common scenario in computer tomography. These approaches often require 3D partial derivatives to be computed, our data simply do not have the requisite (vertical) resolution to facilitate that. Furthermore, as our data are literally mined through manual processes, they only become available progressively at periodic intervals. In our application, slope trends need to be estimated in a causal manner using cross-section samples from as little as two successive benches (at two particular depths). Our problem also differs from grade estimation or material classification for which a variety of probabilistic inference and machine learning approaches are known [12], [29], [32]–[37]. The focus here is explicit spatial representation which encompasses boundary extraction, region correspondence and capturing changes to subsurfaces using differential geometry.

The proposed system takes input in one of two forms. The contours which describe geospatial regions are either given or derived. Fig. 2(a) depicts a common mining scenario where stratigraphic information are obtained by examining core specimen extracted from the drilled holes. Alternatively, geozone transitions (marker bands) are established from geochemical assays or geophysical measurements taken while drilling. The main characteristics of this data is high resolution in z and sparse sampling in the x - y plane. Geologists typically interpolate the boundary at locations (dotted lines) between the drilled holes based on an understanding of the geological setting. The time and effort required to “join up” these vertical slices to form a preliminary 3D surface model

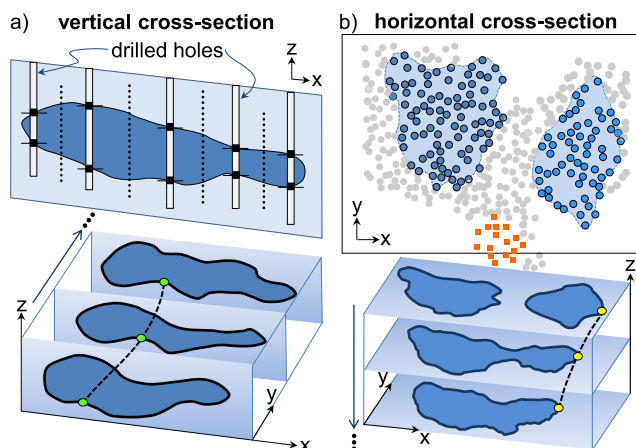


FIGURE 2. Application scenarios. (a) In mining, holes are drilled perpendicular to the strike of the orebody. Geozone transitions (black squares) are commonly established by identifying marker bands using measure-while-drilling or assay data and/or visual inspection of rock specimen. Boundary is interpolated at locations (dotted lines) between the drilled holes. The estimated boundary is subsequently extrapolated and joined in the orthogonal direction, i.e., between successive vertical cross-sections to generate an approximate 3D orebody surface model. (b) Blastholes are more densely and non-uniformly sampled in the x - y plane. In this case, the boundaries shown in the horizontal cross-sections are not actually given. Instead, they are obtained using the proposed boundary extraction techniques. Modeling how the volume evolves through the cross-sections (as hinted by the dash lines) and estimating the slopes are the primary objectives of this work.

can be substantial. Fig 2(b) depicts a second scenario where horizontal cross-sections are taken. This scenario differs from the previous through denser sampling in the x - y plane and having relatively low z -resolution; blastholes are plentiful albeit non-uniformly sampled. In this case, the contours are *not* given — only the coordinates and geozone designations (classification labels) of the blastholes are known. This second scenario poses additional challenges with respect to *boundary extraction* which will be further considered in this paper. These pictures highlight the main objective of this work which is to model how a volumetric region evolves through the cross-sections using contours and estimate the direction in which the boundary moves in.

The proposed boundary geometry modeling framework (henceforth, abbreviated as BGM) consists of four subsystems. The processing pipeline is shown in Fig. 1. The role of each subsystem is briefly described below.

- *Boundary extraction* finds connected regions (components) from labeled blastholes. It locates boundary samples on each component and converts them into a closed contour.
- *Spatial correspondence* automates the process of region association and component alignment. Having identified one or more components on each cross-section, the next goal is to match-up and motion-compensate contours, i.e., find the optimal translation which brings corresponding regions into alignment. By convention, contours found in the top and bottom of any two successive slices are referred as ‘source’ and ‘target’ components respectively.
- *Contour metamorphosis* uses particle trajectories to facilitate slope estimation. The goal is to model residual differences after source–target components are aligned in a common frame. Contour boundaries are embedded in a 2D level-set function, by tracking the movement of the zero-interface, one establishes pathways for morphing the source region(s) into the target region(s) as the level-set evolves under the dynamics of a PDE.
- *Contour prediction* reconstructs a 3D volume through interpolation or extrapolation of normalized trajectories.

C. CONTRIBUTIONS

The main contributions are as follows.

- For boundary extraction, an edge map synthesis procedure is described which converts sparse non-uniform data points to an image representation. Gradient vector field and active contours are used to extract shapes (regularized contours) from unordered edge pixels.
- For spatial correspondence, region association and component alignment problems are solved using FFT cross-correlation under spatial constraints. Obstacle avoidance is approached from a resource allocation view point, multiple-source multiple-target component relationships are explored using intersection graphs.
- For contour metamorphosis, slope estimation is achieved using particle trajectories. A curvelet

TABLE 2. Techniques employed in boundary geometry modeling.

Technique	Field of origin / inspiration
BOUNDARY EXTRACTION	
Connected component analysis	data analysis, clustering
Orientation selective gap closure	mathematical morphology
Synthesis of edge structures	image processing
GVF (anisotropic diffusion)	computational physics
Active contours (object localization)	computer vision
SPATIAL CORRESPONDENCE	
Region association	probability / statistics
Cross-correlation/shift estimation	signal processing (FFT)
Component alignment strategies:	resource allocation,
resolve conflicts, avoid obstacles	perception
Source–target dependency tree	graphical decomposition
Land mass characterization	shape analysis
CONTOUR METAMORPHOSIS / PREDICTION	
Level-sets, signed distance function	applied mathematics
Wave propagation (hyperbolic PDE)	computational physics
Particle trajectories	object tracking
Boundary tracing	topology, computer vision
Surface reconstruction from slices	computer graphics

backtracking algorithm has been devised to overcome tracking failure caused by branching; a phenomenon that leads to boundary distortion. This occurs when there is a significant mismatch in shape (local curvature) between the corresponding boundary segments.

Table 2 provides an overview of the techniques employed in each area. In the ensuing sections, the reasons for choosing these techniques will be elaborated as specific challenges are described. We discuss how standard approaches are modified to address particular needs. The solutions sought for different parts of the problem draw inspiration from different fields.

II. BOUNDARY EXTRACTION

The first objective is to extract the boundary of contiguous spatial regions (connected components) given sparse data. For input, we have the spatial coordinates of blastholes for various horizontal cross-sections, and these blastholes have been classified as belonging to different geozones from prior analysis based on material properties or chemical composition. An example of this is shown in Fig. 2(b) where the blue, gray dots and orange squares represent blastholes sampled from three different geozones. An important point about these samples is that they are irregularly spaced and unordered. Consequently, efficient algorithms for labeling connected components [38], [39] that perform sequential scans on a uniform grid cannot be used here. A related problem is that attempts at forming an edge map by connecting peripheral samples often produce false edges and loops which make clean boundary extraction difficult. These issues will be addressed in due course. In preparation for what follows, the boundary extraction process is summarized by a series of steps.

- **Connected component analysis:** For each cross-section, identify separate clusters within each geozone.
- **Blasthole boundary detection:** Find samples located on the boundary of each cluster / connected component.

- **Edge map synthesis:** Project the region (edges connecting the boundary samples) onto an image grid.
- **Gradient vector field computation:** This is used to drive the active contour (PDE solution) toward the boundary.
- **Active contours evolution:** Extract the boundary as a closed contour described by N control points.

A. CONNECTED COMPONENT ANALYSIS

The modification involves using a kD-tree for nearest neighbor search and adopts a region growing approach. All samples from the same geozone and same cross-section are numbered from 1 to n , initially each belonging to the cluster of ‘self’. Neighboring samples within a radius of r are merged with the current sample and labeled with the minimum cluster index amongst the group. Cluster membership information is propagated iteratively until no further changes occur and S connected components remain.

B. BLASTHOLE BOUNDARY DETECTION

For each connected component, boundary samples are identified by thresholding the local entropy which is significantly non-zero at geozone transition points. Suppose a sample n has N_n neighbors within a radius of r and the fraction of samples belonging to geozones g_1 and g_2 are $p_{n,1}$ and $p_{n,2}$. The local entropy is computed as $h_n = -\sum_i p_{n,i} \log_2(p_{n,i} + \epsilon)$. Sample n is marked as a boundary sample if $h_n \geq \max\{T_{\text{entropy}}, h_n^{(\text{median})}\}$ where $T_{\text{entropy}} = 0.5$ and $h_n^{(\text{median})}$ [the median entropy in n ’s neighborhood] is used to suppress “non-maximum” responses.

C. AUTOMATED GAP CLOSURE

One limitation of this detector is the entropy tends to zero when the geozone labels no longer vary. This creates a problem as boundaries remain essentially open at the frontiers of surveyed regions which are not bordered by a different geozone. To remedy this situation, we perform orientation analysis to close these gaps. The objective is to recognize samples on the outskirts of a geozone as edges. The direction of the K_{orient} closest neighbors from n are computed and sorted in ascending order. If a gap larger than T_{orient} radian is found, sample n is deemed to be on an open edge that needs to be closed. Considering the blastholes are often sampled on a hexagonal lattice, we set $T_{\text{orient}} = \frac{2}{3}\pi$ and $K_{\text{orient}} = (4 \times 2\pi / T_{\text{orient}}) = 12$. Some intermediate results are shown in Fig. 3.

One observation from Fig. 3(b) is that there is often a large gap between [what humans perceive as] adjacent samples along the [as yet undetermined] component boundaries. Indeed, the variable gap size renders useless boundary tracing algorithms like ‘marching squares’ that rely on fixed separating distances. A more robust approach to boundary extraction is to formulate it as an energy functional minimization problem that provides a tradeoff between smoothness and

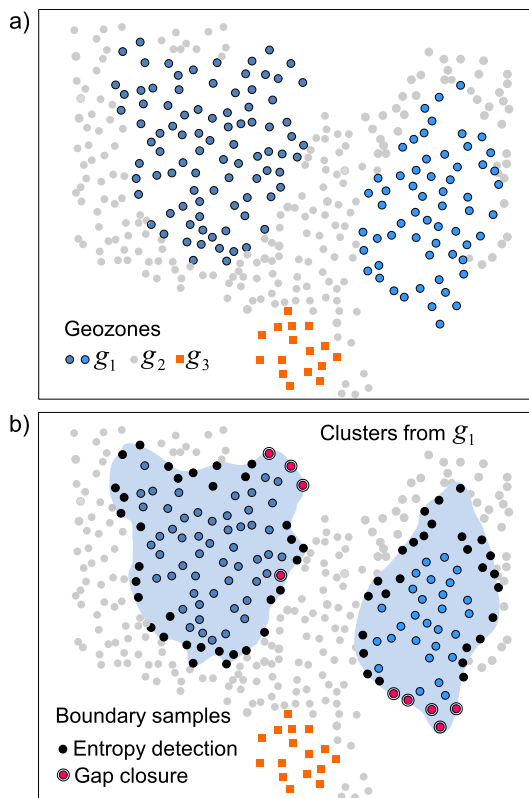


FIGURE 3. Connected component analysis and blasthole boundary detection. (a) Input provides blasthole coordinates and geozone labels, one horizontal cross-section is shown. (b) Two clusters (connected components) have been identified in geozone g_1 ; black dots represent boundary samples detected by thresholding local entropy, red dots represent gap closure informed by orientation analysis.

fidelity. For this, projection of the boundary onto an image grid constitutes the first step.

D. EDGE MAP SYNTHESIS

The main objective is to bridge the gap between adjacent samples located on the boundary. Each sample connects with its K nearest neighbors by forming an edge between them. This edge structure comprising of line segments is then densely sampled and transferred over to the image plane. Specifically, edge energy is accumulated by pixels within some margin of each line segment.³ This procedure is further described in Algorithm 2.⁴ For illustration, a synthesized edge map is shown in Fig. 4(a). Although it contains some false edges, the active-contour segmentation approach (to be described next) can tolerate small imperfection.

E. ACTIVE CONTOUR EVOLUTION IN GRADIENT VECTOR FIELD

The final goal is to obtain a contour (polygon of regularized boundary points) given the synthesized edge structure. This

³Morphological closing is subsequently applied to give the edges adequate thickness.

⁴Mapping real-world coordinates to the image domain (and vice-versa) involves scale changes. Uniform quantization is used implicitly throughout.

is achieved by performing region segmentation using Active Contour Evolution (ACE) in a Gradient Vector Field (GVF). The GVF is obtained from the synthesized edge structure following the procedure given in Appendix A. It may be visualized as a field of arrows (an external force) that pushes any given point in space toward the boundary (see Fig. 4(b)). Accordingly, an active contour which has been initialized as the component bounding box will evolve over time under the action of the GVF and be drawn inward until it converges at the region's boundary. This is illustrated in Fig. 4(c) and the short video (see multimedia content). The computational aspects of GVF active contour evolution are described in Appendix A. This highlights the first connection with computational physics.

III. SPATIAL CORRESPONDENCE

The *boundary extraction* subsystem produces a collection of contours. For instance, a set of real component boundaries extracted from multiple cross-sections can be seen in Fig. 5. When given two successive cross-sections, we will henceforth refer to the upper and lower cross-sections (likewise for the components contained within them) as 'source' and 'target' respectively. This section presents general strategies for establishing relationships between source and target contours (viz., *region association*) and estimating the optimal translation (performing *component alignment*) with and without spatial constraints. The devised strategies take into account unique characteristics such as incompleteness, causality and resolution disparity in the data as foreshadowed in the introduction.

In this paper, spatial correspondence is treated as a two-stage process. The spatial mapping is the result of combining component translation and contour metamorphosis (local surface deformation). The present goal is to account for the rigid movement of associated components — principally the translation observed in successive cross-sections. Subsequently, non-rigid movement will be considered in Section IV.

A. REGION ASSOCIATION

Given two cross-sections with n_s source components and n_t target components, the goal is compute the association matrix $A \in \mathbb{R}^{n_s \times n_t}$ which describes the relationships between s and t . In the preliminary assessment, a '1' in $A(s, t)$ implies there is potential interaction between source s and target region t .

1) FORMULATION

A Gaussian pdf $p_s(\mathbf{x} \mid \mathbf{d}) \propto \exp(-a(\|\mathbf{x} - \mathbf{c}_s\|/b)^2)$ is used to model the likelihood of source-target association. In this setup, spatial proximity is measured by the difference of \mathbf{x} : an arbitrary location in the image and a fixed \mathbf{c}_s : the centroid position of source component s . Design choices revolve around the hyper-parameters a and b which must capture relevant aspects to produce sensible results. The right

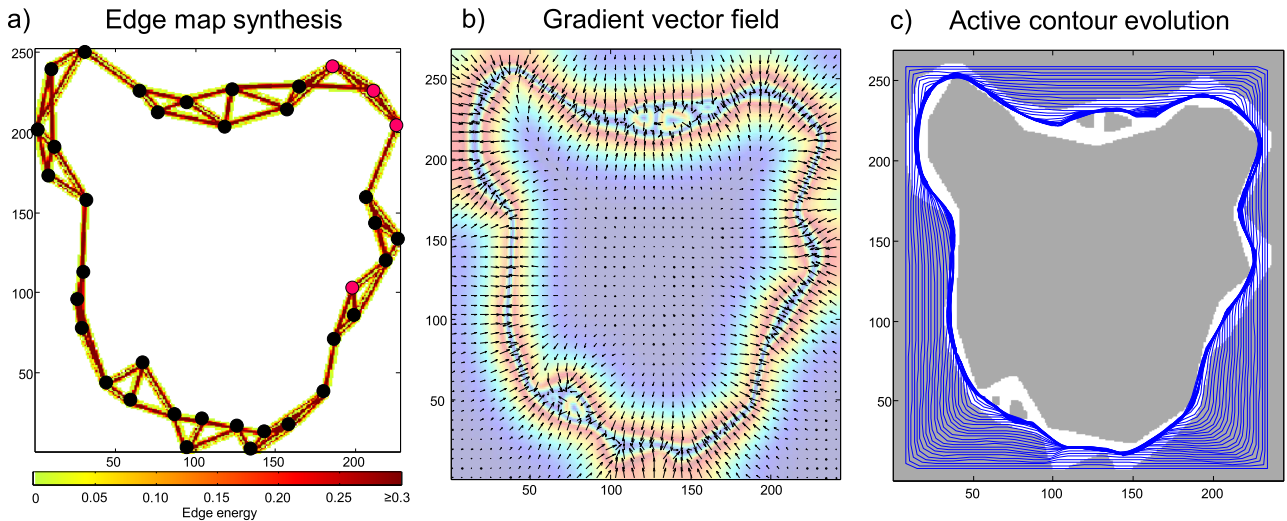


FIGURE 4. (a) Synthesized edge map for a connected component after step 5 in Algorithm 2, boundary blastholes are overlaid. (b) Gradient vector field computed from the edge map. (c) Active contour evolves in the GVF, gravitating toward the region boundary at steady state. The default parameters are $K_{struct} = 5$, $K_{joints} = 4$, $n_x \times n_y \approx 240^2$, $p = 8$.

formulation very much depends on the problem, the sampling characteristics and reliability of the data.⁵

Apparent motion and *object area* represent two factors most relevant to our model. Apparent motion is inferred from the statistics $\mu_{\mathbf{d}_{min}}$ and $\sigma_{\mathbf{d}_{min}}$ which are computed from observed distances between the sources and targets, $d_{s,t}$ as follows.

$$\mathbf{d}_{min} = \left[\min_t d_{s,t} \right]_s \in \mathbb{R}^{n_s \times 1}, \quad (1)$$

$$\mu_{\mathbf{d}_{min}} = \mathbb{E}[\mathbf{d}_{min}], \quad \sigma_{\mathbf{d}_{min}} = \sqrt{\mathbb{E}[\mathbf{d}_{min}^2] - \mu_{\mathbf{d}_{min}}^2} \quad (2)$$

Variability in the observations is given by the standard error $s_{\mathbf{d}_{min}} = \sigma_{\mathbf{d}_{min}} / \sqrt{n_s}$. The association probability takes the form

$$p_s(\mathbf{x} | \mathbf{d}_{min}) \propto \exp\left(-\frac{(\|\mathbf{x} - \mathbf{c}_s\| / \mu_{\mathbf{d}_{min}})^2}{2\nu^2}\right), \quad (3)$$

$$\nu = \frac{4}{1 + \exp(-s_{\mathbf{d}_{min}}/s_0)} \quad (4)$$

The term $s_{\mathbf{d}_{min}}/s_0$ relates to the uncertainty of the observations and is normalized by a reference value s_0 . The ‘noise’ parameter $\nu \in [2, 4]$ regulates the shape of the Gaussian. When variability is low ($s_{\mathbf{d}_{min}}/s_0 \rightarrow 0$) the association function contracts and becomes more concentrated around \mathbf{c}_s . Conversely, it dilates when there is a lack of consensus on

⁵Examples of probabilistic reasoning and state estimation techniques are plentiful in the literature. Those that incorporate geometry information for visual tracking [40]–[46] generally rely on having rich and accurate information delivered at a high frame rate. e.g., solving spatial correspondence problems using optical flow or persistent templates. In our application, shape invariance is not something that can be reliably exploited. Topological changes, under-sampling (large spacing between successive cross-sections) and sampling uncertainty (mixture of material and poor resolution) often contribute to sudden and unexpected changes. In geology and mining exploration, operational constraints also limit the amount of observations that can be gathered, this creates a situation where fast changing dynamics may surpass the learning rate for a complex model. Given the few samples available, we resort to using the most basic measures like the magnitude, but not the directionality, of motion.

$\mathbb{E}[\mathbf{d}_{min}]$; this allows more distant targets to become potentially associated with the source.

2) SCALE EFFECTS

Since the \mathbf{d}_{min} observations are noisy, component size should be compared with the apparent translation $\mu_{\mathbf{d}_{min}}$. When $\mu_{\mathbf{d}_{min}}$ exceeds the object span φ_s , the policy $\mu_{\mathbf{d}_{min}} \leftarrow \min\{\varphi_s, \mu_{\mathbf{d}_{min}}\}$ should be adopted. This limits the scope of association for smaller components by shrinking $p_s(\mathbf{x} | \mathbf{d}_{min})$.⁶ To prevent $p_s(\mathbf{x} | \mathbf{d}_{min})$ from becoming too narrow — which would eliminate any source-target association potential when taken to the extreme — it would be prudent to also place a lower limit (μ_{lower}) on $\mu_{\mathbf{d}_{min}}$.

The final form of the association likelihood is given by

$$p_s(\mathbf{x} | \mathbf{d}_{min}) \propto \exp\left(-\frac{(\|\mathbf{x} - \mathbf{c}_s\| / \lambda_s)^2}{2\nu^2}\right), \quad \mathbf{x}, \mathbf{c}_s \in \mathbb{R}^2 \quad (5)$$

$$\lambda_s = \max\{\min\{\mu_{\mathbf{d}_{min}}, \varphi_s\}, \mu_{lower}\} \quad (6)$$

$$\nu = \frac{4}{1 + \exp(-s_{\mathbf{d}_{min}}/\mu_{lower})} \quad (7)$$

where μ_{lower} is related to the anticipated lateral movement. It basically dictates the minimum radius of association.

3) DECISION FUNCTION

The decision of whether to associate source s with target t is governed by a discriminant function

$$A(s, t) = \begin{cases} 1 & \text{if percentile}(\{p_s(\mathbf{x} | \mathbf{d}_{min}) | \mathbf{x} \in \mathcal{R}_t\}, K) \geq 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

$$K = 50 \times \left(1 + \left(1 - \min\left\{\frac{a_s}{a_t}, 1\right\}\right)\right) \in [50, 100] \quad (9)$$

⁶The object span is approximated from the area as $\varphi_s \approx \sqrt{a_s/\pi}$.

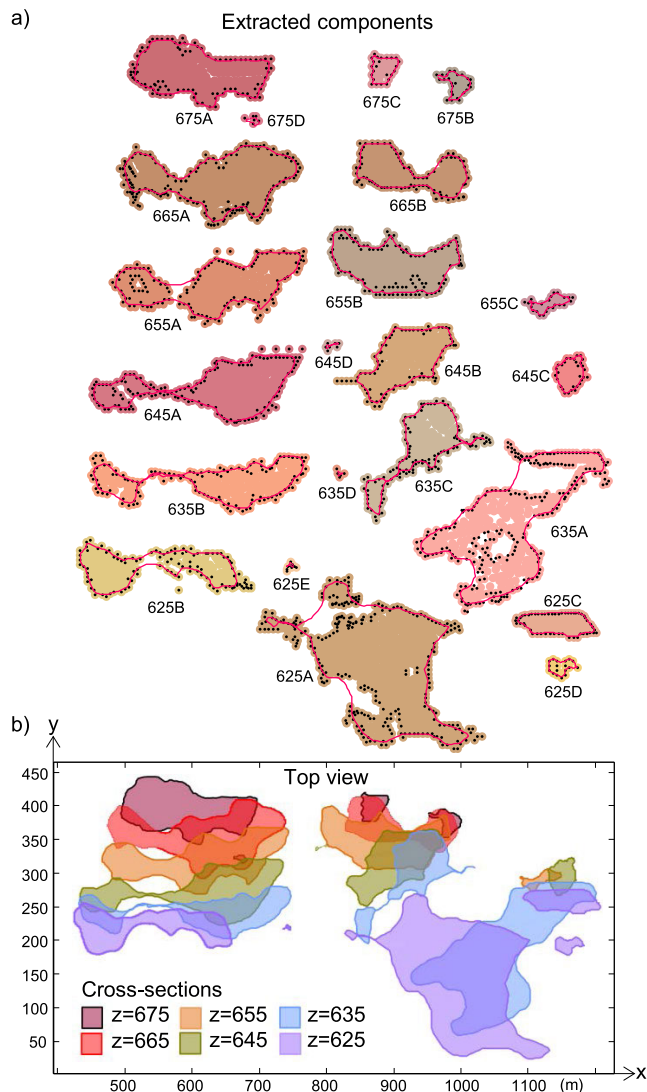


FIGURE 5. Component boundaries extracted using active contours for six horizontal cross-sections (spaced 10m apart) in one geozone. Black dots denote blasthole samples located at / near geozone transitions. These serve as input to the next subsystem in the boundary geometry modeler called ‘spatial correspondence’. The goal is to associate regions and estimate component alignment between corresponding source and target contours subject to spatial constraints.

- The set notation $\{p_s(\mathbf{x} \mid \mathbf{d}_{\min}) \mid \mathbf{x} \in \mathcal{R}_t\}$ in (8) refers only to points inside the support interval of target component t .
- The percentile K in (9) depends on the source-to-target area ratio (a_s/a_t). The rationale is as follows:
 - If $a_s < a_t$, the source can (at best) fit itself wholly within the target with room to spare. Thus, a fair value for assessing spatial proximity (the association potential) is the median likelihood, offset by the fraction of a_t which s cannot overlap.
 - In the limit as $(a_s/a_t) \rightarrow 0$, K seeks out the maximum likelihood when s is treated as a point object.

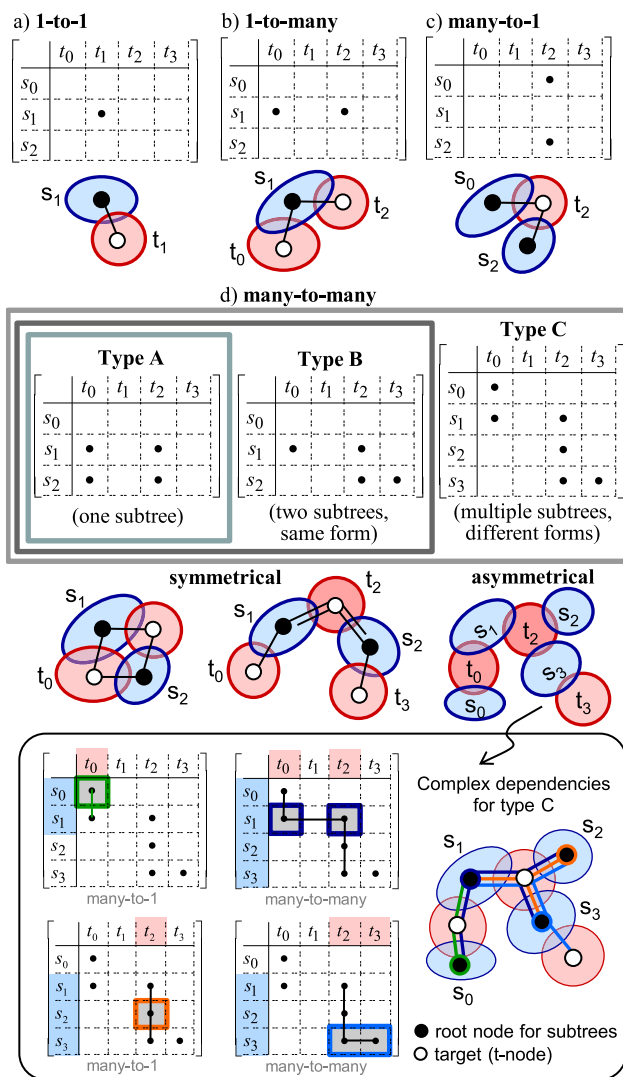


FIGURE 6. Source–target component relationships (viz., pattern of potential overlap) are represented by intersection graphs given the association matrix. Source and target components are denoted by black and white nodes respectively. Different region association scenarios such as multiple-source single-target are depicted in (a)–(d). For complex scenarios (see inset), suitable component alignment strategies will be applied to each subtree anchored at an s-node which includes all connected t-nodes and their children s-nodes.

B. COMPONENT ALIGNMENT STRATEGIES

Having obtained the association matrix $A(s, t)$, source–target relations can be represented by a graph. This gives rise to several possibilities: a 1-to-1 mapping, many-to-one, one-to-many, and many-to-many correspondence between the source and target components. These scenarios are shown in Fig. 6. For complex scenarios (e.g., multiple sources associated with multiple targets), the intersection graph is decomposed into subtrees (see Fig. 6 inset). This section is concerned with devising suitable strategies for each scenario, i.e., finding source translation estimates which maximize target alignment or component coverage. Graphically, each scenario is described by a subtree rooted at an s-node and its expansion

includes all connected t-nodes and their children (any s-nodes connected to the t-nodes).

1) DISPLACEMENT ESTIMATION: CROSS-CORRELATION VIA FFT

A 1-to-1 mapping, or single-source single-target alignment, is the simplest situation encountered in region association. Conceptually, the best alignment is achieved by template matching; this is otherwise known as 2D cross-correlation. Using vectorized notation, e.g., $\mathbf{i} = (i_0, i_1)$, the problem may be posed as

$$\text{optimal translation} = \underset{\mathbf{m}}{\operatorname{argmax}} r_{VU}[\mathbf{m}] \quad (10)$$

where the matched filter response with lag \mathbf{m} is given by

$$r_{VU}[\mathbf{m}] = \sum_{\mathbf{i}} v[\mathbf{i}] u[\mathbf{i} + \mathbf{m}] = \sum_{i_0, i_1} v[i_0, i_1] u[i_0 + m_0, i_1 + m_1] \quad (11)$$

These operations can be implemented efficiently as

$$r_{VU}[\mathbf{m}] = \mathcal{F}^{-1}(V[\mathbf{k}] U^*[\mathbf{k}]) \quad (12)$$

using the Fast Fourier Transform (separable FFT) [47]. In this equation,⁷ r_{VU} represents the 2D cross-correlation between u and v , where $u \equiv \mathcal{U}$ is the source mask (a binary indicator function for the region occupied by the source component in the image plane) and $v \equiv \mathcal{V}$ is a signed distance function [48] (SDF) for the target component. The SDF (computed in Algorithm 3) gives a 2D map of the signed distance of any location from the contour boundary (zero-interface) and adheres to the convention where pixels inside and outside the target region have positive and negative values, respectively.

On the RHS, \mathcal{F}^{-1} represents the inverse FFT. $U^*[\mathbf{k}]$ and $V[\mathbf{k}]$ denote the complex conjugate of the FFT coefficients for u and FFT coefficients for v . Finding the optimal shift (displacement \mathbf{m}) is equivalent to locating the peak in r_{VU} . This lays the basic foundation for component alignment (see Algorithm 4).

2) MULTIPLE-SOURCE SINGLE-TARGET SCENARIO

The proposed solution (10) requires certain modifications in a multiple-source single-target (many-to-one) scenario. Consider the situation depicted in Fig. 7(a). As it stands, if source s_2 acts selfishly and tries to maximize its correlation with the target with no regard for other source components, s_1, s_3 and s_4 will be locked out of the region merging arrangement. To prevent this, the sources will participate in target sharing. This is posed as a *resource allocation* problem where the objective is to divide the target into sectors (or ports of varying sizes) and assign the ports in such a way as to minimize their angular difference with the affiliated source component (see Fig. 7(a)). The solution is presented in Algorithms 5–6.

⁷Caaveats: Certain technical conditions need to be satisfied for this to work. Stated simply, the inputs need to have first quadrant spatial support and be sufficiently padded around the margins to ensure no shifted component ever steps outside the image border. Detailed guidance is given in [47] §4.6.4.

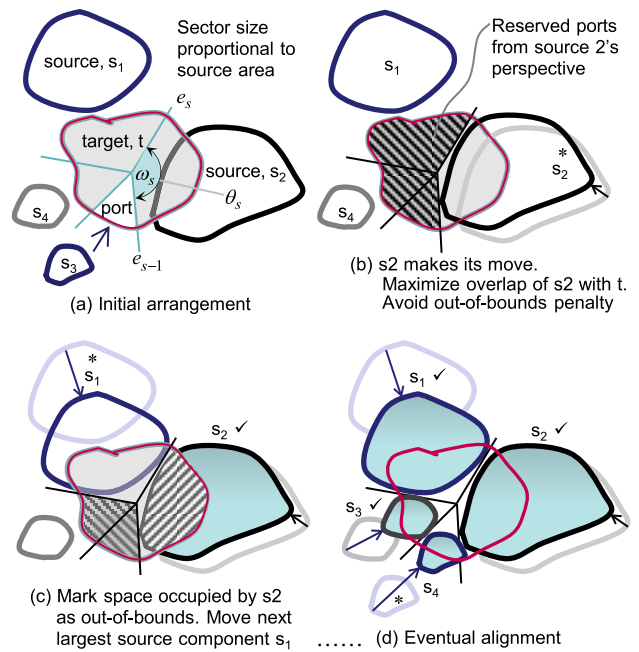


FIGURE 7. Port allocation concept for multiple-source single-target scenario.

The port boundaries and target sector-to-source mapping so obtained allow very specific spatial constraints to be imposed. For instance, in Fig. 7(b), by imposing a penalty on target sectors reserved for other source components (see striped region), s_2 will refrain from stepping into other’s territory whilst maximizing its overlap with the target. This behavior can be realized by setting $v[\mathbf{i}]$ (initially, the target signed distance function) to large negative values at locations $\mathbf{i} \equiv (i_0, i_1)$ which are marked out-of-bounds.

3) SINGLE-SOURCE MULTIPLE-TARGET SCENARIO

An issue with the current approach is that it is somewhat biased toward single-target alignment. There is no *incentive* for a source component to overlap with multiple targets since stepping outside the boundary of a target incurs a negative penalty in (11). This usually deters any ‘cross-over’ unless the source component is large enough to encompass both targets and the targets are reasonably close together that reward outweighs penalty. As motivation, Fig. 8(b) illustrates a situation where the source is encouraged to straddle two targets; this in turn maximizes component coverage and offers a more plausible explanation in a region splitting scenario. Within the current framework, this can be achieved by placing rewards at strategic locations in the value function v to promote cross-over. This is further described in Algorithm 7.

4) MULTIPLE-SOURCE MULTIPLE-TARGET SCENARIO

The *cooperative* and *reward* strategies described in §III-B2–III-B3 provide a unified way for dealing with complex relationships under the displacement estimation framework of §III-B1. As an illustration, these techniques are

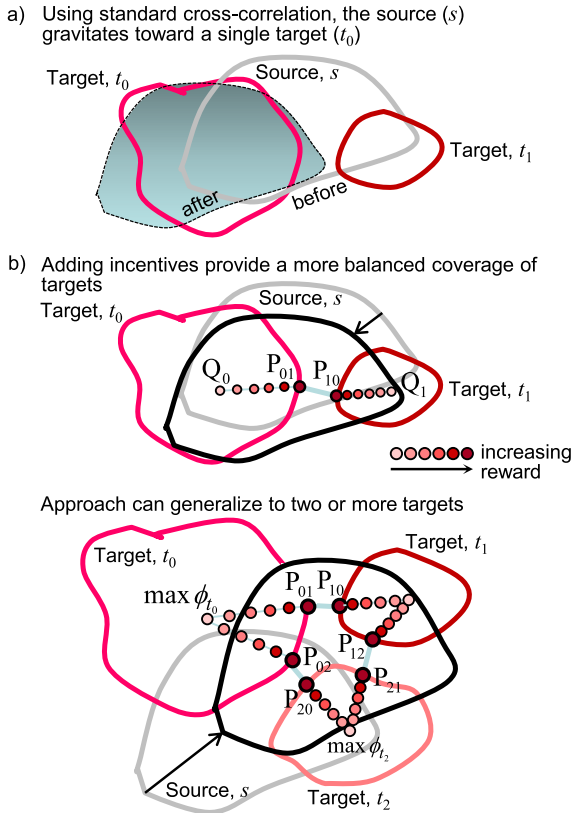


FIGURE 8. Target coverage in single-source multiple-target scenario.

applied to a multiple-source multiple-target scenario in Fig. 9. The initial position of the source and target components are shown in Fig. 9(a). This has the same graphical structure as the type C configuration in Fig. 6. Hence, it can be decomposed into 4 subtrees or expansion steps as shown in Fig. 9(b). A compact way of writing these steps is ①: $\mathbf{s}_1 \rightarrow \{t_0 \leftarrow s_0, t_2 \leftarrow (s_2, s_3)\}$, ②: $\mathbf{s}_0 \rightarrow t_0 \leftarrow s_1$, ③: $\mathbf{s}_2 \rightarrow t_2 \leftarrow (s_1, s_3)$, ④: $\mathbf{s}_3 \rightarrow \{t_2 \leftarrow (s_1, s_2), t_3\}$ where the source component under consideration (to be moved) is typed in bold.

Within a subtree, each link describes a tentative association between an s-node and a t-node in Fig. 9(b). This simply means there is a potential for the source to intersect with the target. The objective is to eliminate edges where a connection does not exist — perhaps due to obstacles (the presence of other source components) which have prevented an overlap between s and t — and find the translation for source components that do overlap with targets (i.e., maximize correlation subject to spatial constraints).

For brevity, the subtree expansion steps are omitted here. Interested readers may refer to the worked example given in the Supplemental Material for a full description of the expansion steps. Through the lens of spatial reasoning (Algorithm 3–9 and Figure 4–9), Section III has outlined the main connections with signal processing, computer vision, resource allocation and graphical decomposition in this work.

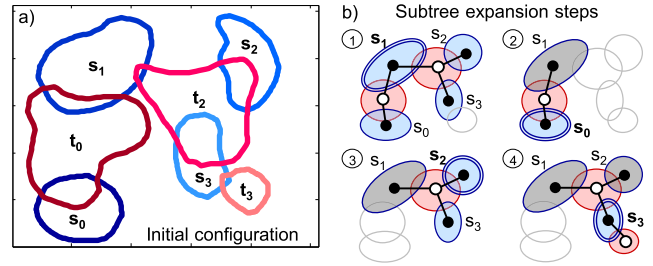


FIGURE 9. Spatial correspondence in multiple-source multi-target scenario. Refer to the Supplemental Material for a worked example with details of the expansion steps.

In summary, the *spatial correspondence* subsystem produces an association matrix A and translation estimates \mathcal{T} which describe the relationship between source and target components in two successive cross-sections. Specifically, the source contours $\{C_s\}$ associated with a given target C_t define an instance for the next subsystem, *contour metamorphosis*, to work on.

IV. CONTOUR METAMORPHOSIS

The objective is to capture local spatial variation using differential geometry, to explain how source regions evolve into target regions using PDE. This problem may be visualized as a transformation from (a) to (b) in Fig. 10.

Given displacement estimates $\{\mathbf{m}_s = (m_s^{(x)}, m_s^{(y)})\}_s$ and translation operator $\mathcal{T}(\mathbf{x}, \mathbf{y}) = (\mathbf{x} + m_s^{(x)}, \mathbf{y} + m_s^{(y)})$, the starting point is a principal target contour $C_t \equiv (\mathbf{x}_t, \mathbf{y}_t)$, its associated (shift-compensated) source contours $\{C'_s \equiv \mathcal{T}(\mathbf{x}_s, \mathbf{y}_s)\}$ and any affiliated target contours $\{C_{t'}\}$ connected to the source nodes. Fig. 10(a) shows a simple case where there is only one source and one target contour involved. It is this residual difference, the non-rigid changes in shape in the aligned common frame, that is being modeled.

In the PDE model, contours are represented by level-sets ϕ_t which embed region boundaries as the zero-interface, viz., $\Omega_t = \{\mathbf{n} \in \mathbb{R}^2 \mid |\phi_t[\mathbf{n}]] \leq 0.5\}$. Initially, $\phi_{t=0}$ is set to γ_s , the signed distance function of C_s . The basic objective is to model the evolutionary process from $\phi_{t=0} = \gamma_s$ to $\phi_{t=\infty} = \gamma_t$. As the level-set $\phi(t)$ evolves,⁸ the region boundary morphs from (a) to (b). The morphing process is described in Appendix B wherein (20) represents the key equation for level-set update. Visually, this computation describes a *propagating wavefront* as depicted in Fig. 10(c).

A. ISO-CONTOURS AND WAVE SPEED ADJUSTMENT

It is important to note that these level-set zero interfaces Ω_t are functions of time. Ultimately, we need to obtain iso-contours at different elevations z . However, no single Ω_t (other than Ω_0) currently corresponds to a fixed elevation.

As a corollary, not all points on the advancing wavefront will reach the target contour at the same time since the wave

⁸A level-set may be defined as a collection of iso-contours or a family of embedded point sets each represented by $\mathcal{S}_k = \{\mathbf{n} \mid \phi(\mathbf{n}) \geq k\}$.

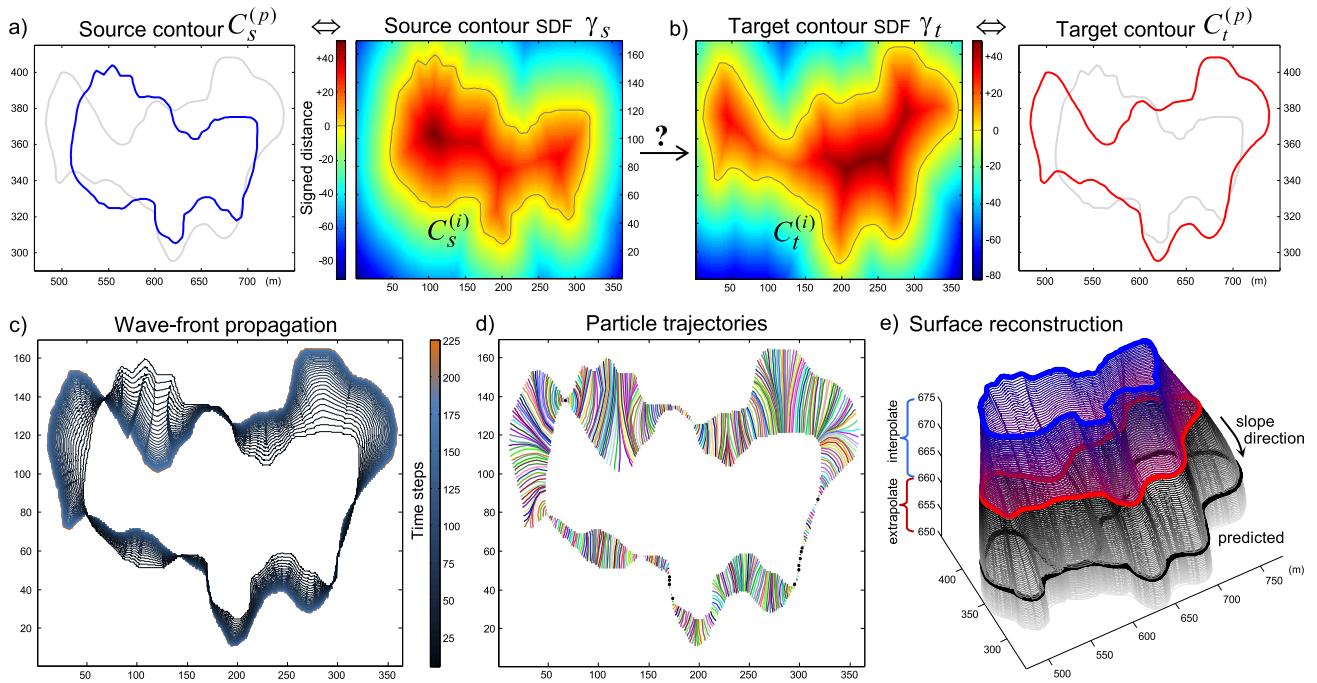


FIGURE 10. Overview of contour metamorphosis. Basic objective is to model how the source region morphs into the target region using differential geometry, viz., explain how (a) is transformed into (b) using the narrative of (c) and (d) to ultimately produce a surface model similar to (e). Notation-wise, superscripts (p) and (i) denote the physical domain and image processing (modeling) domain respectively. Physical coordinates, e.g. $C_s^{(p)} = (x_s^{(p)}, y_s^{(p)})$, are measured true-to-scale and expressed in UTM coordinates in [m]. (a) A source contour (blue line: $C_s^{(p)}$) is aligned with the target contour (light gray) in a common frame, it has an equivalent level-set representation γ_s where contour boundary is embedded as the zero-interface. (b) Associated target contour (red line: $C_t^{(p)}$) and its level-set (signed distance) function γ_t . The arrow annotated with “?” is elaborated in (c) which explains the evolution process as a propagating wavefront, viz., the movement of the zero-interface as (a) morphs into (b). In terms of tracking, (d) provides an alternative description using particle trajectories which emphasizes on directionality.

propagates at different speeds depending on location.⁹ Therefore, *particle trajectories* (see Fig. 10(d)) are used to track the direction in which the zero-interface moves to facilitate slope estimation at a specific height. The relevant procedure called “particle advection” is described in Appendix B wherein (26) represents the key equation for particle update.

The particle speed (spacing) along each trajectory will eventually be adjusted to ensure each advances at a uniform rate and reaches the target boundary at the same time. This allows the zero interfaces to be converted into iso-contours of constant elevation. These ideas capture the essence of contour metamorphosis; Fig. 10(e) also serves as a prelude of the end goal.

B. CHALLENGES

In principle, Appendix B provides all the necessary tools for capturing information on local spatial deformation. In practice, when the shape of the source and target regions are dissimilar, the scheme may fail. Specifically, particle trajectories may fail to track the wavefront in areas where complex topological changes occur especially when they coincide with large differences in local curvature. This can be seen in Fig. 11(a) where some target boundary segments have

⁹The propagation speed depends on the geometric distance between the interface and target contour in accordance with (19).

zero particle coverage. Fig. 11(b) highlights regions which end up being unmodeled despite being reached by the wavefront. This usually occurs due to undersampling when the cross-section spacing is large. The goal in this section is to determine its root cause and contribute a solution to this problem.

In each time step, wavefront segments which the particle trajectories had failed to track are archived. This produces a collection of time-stamped pixels (called **curvelets**) which identify areas of innovation, viz., locations where the wavefront and head of the particle trajectories have diverged. The window of interest in Fig. 11(c) is magnified in (d) to reveal its structure. It can be seen the waves emanate from two separate sources and merge in the middle, this describes a region splitting situation which can be reasonably expected based on the topology shown in (a). For other unmodeled regions shown in (c), the situation is less complex and the behavior can be attributed to significant differences in local curvature. The common theme is that the missing particle trajectories appear to emerge from a single **branching point**.

The color strands in Fig. 11(e) show the trajectory flow for new particles sitting on the curvelets. The arrows indicate the direction the **backtracking** algorithm proceeds in. The aim is to establish pathways from the target back toward the branching point. In Fig. 11(f), new pathways deduced

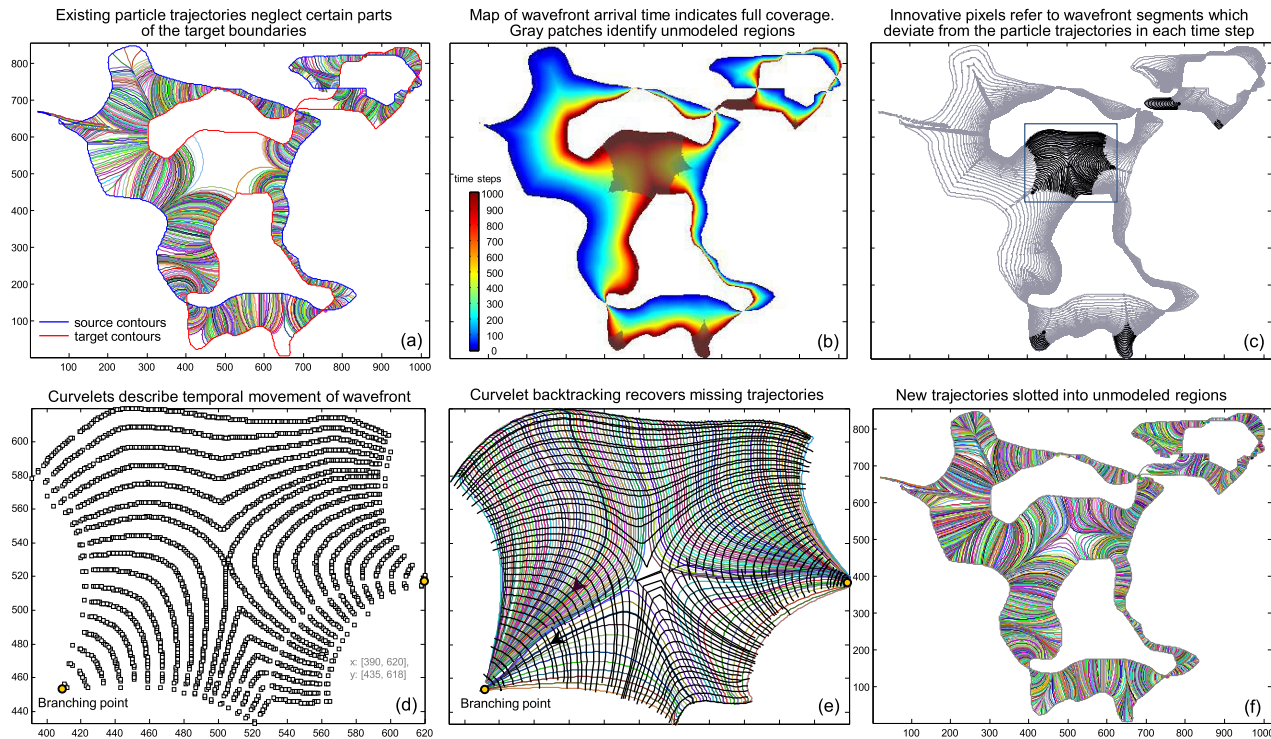


FIGURE 11. Illustration of branching point phenomenon. (a) Particle trajectories (obtained using Appendix B) reveal target segments with zero particle coverage (see empty red lines). (b) Time of arrival map indicates the target segments concerned were clearly reached by the wavefront, gray area indicates region with no trajectory coverage. (c) Innovative areas where the wavefront diverges from the evolving particle trajectories are identified by black lines in each time step. (d) Magnified view of the window in (c) demonstrating the concept of curvelets, a collection of pixels (only a subset is shown), which describe the temporal movement of the wavefront, specifically, the portions which the particles have lost track of. These emanate from a single branching point. (e) Curvelet back-tracking between successive time steps as a means of re-establishing the missing trajectories. Final result: (f) Merging the recovered curvelet trajectories with existing particle trajectories.

from the curvelets are slotted into the unmodeled regions. Evidently, this recovers directional information between the source and target contours in areas where information had previously been lost.

C. CURVELET BACKTRACKING

The backtracking algorithm is formally described in Algorithm 10. The starting point is a collection of time-stamped pixels, these represent wavefront segments the particle trajectories had failed to track using the particle update equation (26) during contour metamorphosis. Backtracking comprises six primitive steps; these are described below and illustrated in Fig. 12.

- (a) **Curvelet endpoint identification** Locate the start and endpoints of prospective curvelets within \mathcal{P}^t — the set of unordered pixels on the wavefront which have drifted away from the particle trajectories at time t .
- (b) **Curvelet segmentation (clustering and ordering)** Extract individual curvelets ζ_c^t by tracing the connecting pixels between relevant endpoints.
- (c) **Curvelet regularization** Fit pixels with a smoothing spline to remove jitter. Resample segment as N_c point regularized curvelet κ_c^t .

- (d) **Curvelet track initialization/extension** The displacement of curvelet points between successive time steps $[\tau_c^t(n), \tau_c^{t-1}(n), \dots]$ describe possible pathways for the missing trajectories in unmodeled regions.
- (e) **Curvelet association** Given regularized points $\kappa_c^{t+1}(m)$ for curvelet c at time $t + 1$, find a matching curvelet k at time t .¹⁰ A case of interest is depicted in Fig. 12(e) where the segmented pixels $\zeta_k^t(n)$ from two curvelets $k = A, B$ are merged from t to $t + 1$. This is handled by considering interval mappings \mathbf{k}^{t+1} for multiple curvelet parts ($p = 1, 2$).
- (f) **Curvelet correspondence** Once matching pixels ζ_k^t are found, the curvelet at time t is smoothed and resampled evenly with N_c points. This curvelet κ_k^t is then collec-

¹⁰Nearest neighbor search is employed to find a surjective mapping from t to $t + 1$. Each point $\kappa_c^{t+1}(m)$ has at least one match at t , but not every pixel on ζ_c^t is matched with a point at $t + 1$. For this reason, as illustrated in Fig. 12(e), this is used only for curvelet association and not for point-wise correspondence. This is because inconsistent (jittery) movement has a detrimental effect on trajectory estimation. To eliminate the cumulative effects of drifting which would otherwise skew the tracks over time, the associated point set is constrained to lie on a regularized curvelet as shown in Fig. 12(f).




FIGURE 12. Curvlet backtracking concepts (cf. Algorithm 10).

tively mapped onto κ_c^{t+1} , the current curvlet.¹¹ For a given point n , the sequence $[\kappa_c^{t+1}(n), \kappa_c^t(n), \kappa_c^{t-1}(n), \dots]$ subsequently defines a track which traverses the unmodeled region.

These steps capture the essence of the backtracking algorithm which produces particle tracks, see color strands in Fig. 11(e) and (f), for the unmodeled regions. The remaining tasks involve branching point identification which enables existing particle trajectories [the trunk portion extending from the source contour to the branching point] to be fused with the newly discovered tracks [from the branching point to the target contour] to complete the missing pieces. The details are described in Algorithms 11 and 12 and the notations are clarified in Fig. 13.

V. RESULTS

The proposed boundary extraction technique was applied to blasthole patterns from a Pilbara iron ore mine in Western

¹¹This provides an interesting contrast to the wavefront tracking approach proposed by Tomek et al. [49] which represents two sets of time-stamped pixels with a complete oriented bipartite graph, poses the task of finding the best tracking arrows as a network flow problem, formulates and solves it as an integer program using branch and bound. Similarly, Singh et al. treated this as an optimization problem and employed dynamic programming in [50].

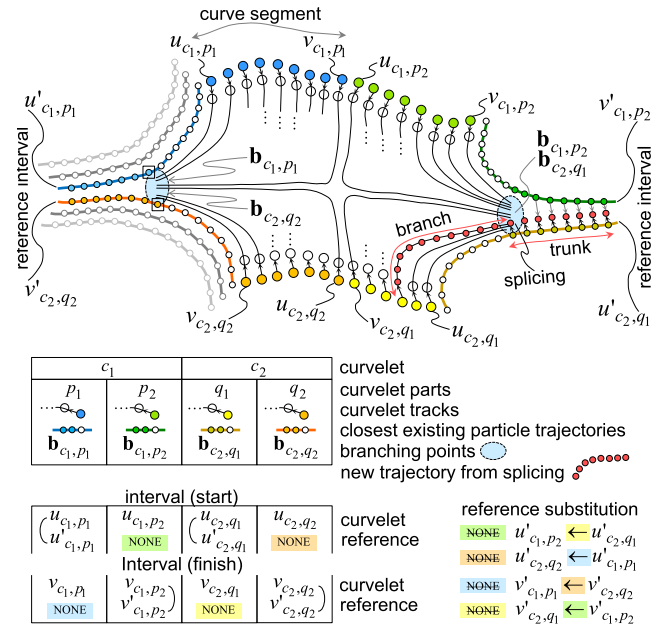


FIGURE 13. Curvlet branching considerations (cf. Algorithm 11–12).

Australia. Fig. 14(a) illustrates one such pattern for a mineralized geozone before any modeling was done.

A. VISUALIZATION

Spatial correspondence and metamorphosis were applied to the segmented regions; culminating in the contours seen in Fig. 14 (b and e). These contours are triangulated to produce the surfaces shown in Fig. 14 (c and f). These results demonstrate that differential geometry can model subterranean boundaries, handling topological changes (e.g., region merging) satisfactorily. The proposed techniques essentially transform irregularly-spaced points of low resolution into continuously deformable surfaces that represent boundaries. Fig. 15 reveals the structure of the two geozones. The two vantage points show that g_1 lies above g_2 , such bedded planes (layered geological formations) are common for iron ore deposits at the mine site.

B. PREDICTIVE VALUE

To assess the utility of slope information conveyed by the model, precision and recall rates are computed for predicted contours and compared against ground truth boundaries. The terrain extends from 690m down to 610m and is partitioned into disjoint 10m intervals (called benches) and processed top-down. Predictions are made using only information above the current bench floor. Therefore, predicted contours represent extrapolations from the known interval. For each comparison, data at the predicted depths from the bench below are withheld (not used to inform the prediction model) and these contours are used only for validation purpose at a given bench. Precision and recall rates are defined using set notations. The general setup is shown in Fig. 16. For predicted region P_i , precision is defined

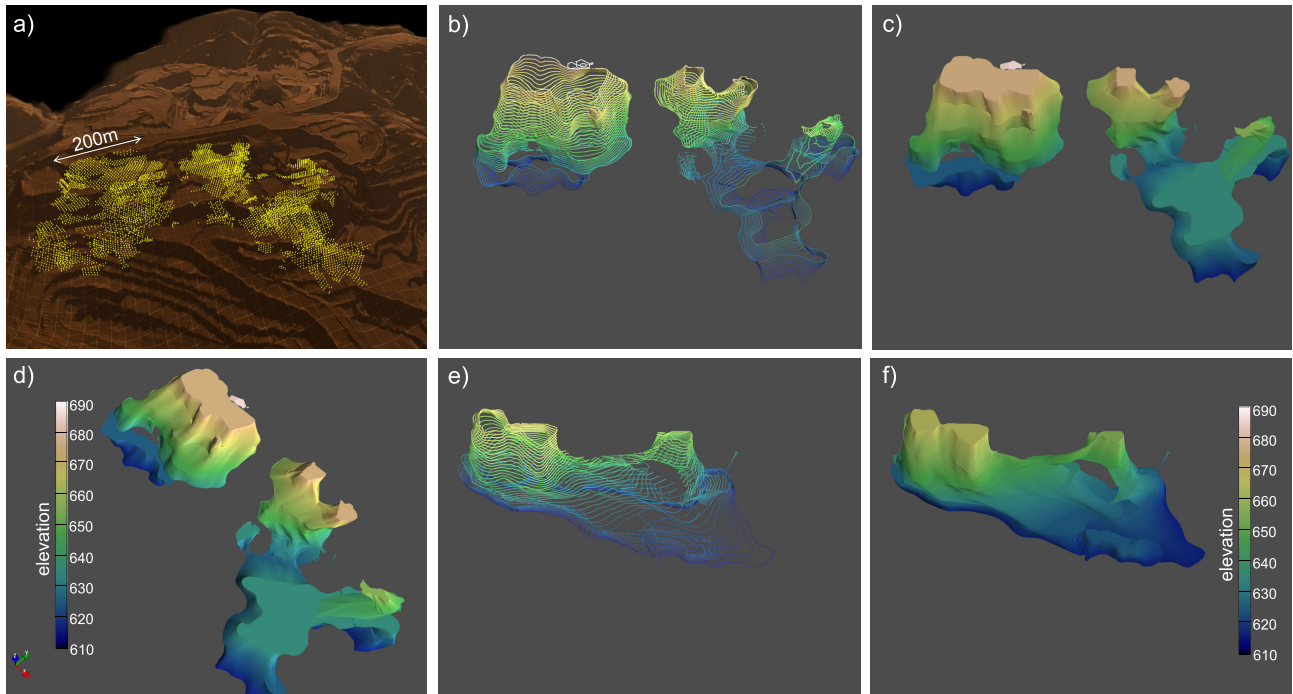


FIGURE 14. (a) Blasthole pattern superimposed on the terrain. Extracted contours and triangulated surfaces for two geozones: g_1 in (b)-(c) and g_2 in (e)-(f). The surface appearing in (c) is shown from a different vantage point in (d).

TABLE 3. Model precision and recall rates vs prediction depth.

Depth (m)	Precision (%)			Recall (%)		
	nil	model	gain	nil	model	gain
1.25	91.0	95.0	+4.0	90.4	97.6	+7.2
2.5	84.6	90.9	+6.3	83.4	95.1	+11.7
3.75	78.7	88.1	+9.4	75.4	90.9	+15.5
5.0	72.0	84.7	+12.7	69.8	87.4	+17.6
6.25	66.5	81.4	+14.9	64.1	82.3	+18.2
7.5	59.8	77.5	+17.7	58.0	77.3	+19.3
8.75	53.8	73.0	+19.2	51.9	73.4	+21.5
10.0	49.1	67.5	+18.4	45.1	67.9	+22.8

as $p(i) = \frac{\sum_j |G_j \cap P_i| |P_i \cap G_j|}{|P_i|}$. For ground truth region G_i , recall is defined as $r(i) = \frac{\sum_k |G_i \cap P_k| |G_i \cap P_k|}{|G_i|}$. The overall precision (resp., recall rate) at depth d is the area-weighted average precision (resp., recall rate) of all contours at the specified depth. The depth varies from 0 to 10m in steps of 1.25m. Precision and recall statistics are computed under two conditions: (i) not using slope information [prediction employs zero-order hold], (ii) using the gradient field from the proposed model to obtain displacement estimates, i.e., combining the translation with the local deformation component estimated during spatial correspondence and contour metamorphosis, respectively.

Fig. 17 demonstrates a consistent improvement in precision and recall rate when slope information is used in the model prediction. Significant gains in precision (12.7–19.2%) and recall (17.6–22.8%) are observed at depths of 5 to 10 meter. These statistics are summarized in Table 3.

VI. DISCUSSION

These results demonstrate the feasibility of modeling subsurface boundary geometry using interdisciplinary techniques. In spite of the incomplete, sparse, causal and limited nature of the spatial patterns, computational physics (partial differential equations) and computer vision techniques can be combined to estimate directionality and produce a subsurface model that gives greater insight into the general structure of a geological domain. Modeling geological boundaries using differential geometry also offers other advantages. For ore reserve estimation, the volume of the triangulated surfaces may be computed easily from the vertices as outlined by Zhang [51]. Given coherent slope estimates, directional predictions made by the model may be used in adaptive sampling [52] to make drilling more targeted and cost-effective. Indeed, various forms of analysis performed on those samples (including chemical assays and material classification by experts) can further constrain or correct inaccurate predictions, when geological boundaries need to be updated or learned further afield. This may be appreciated as a form of active learning in an exploration/exploitation framework. The surface slopes can also serve as a decision support tool for multi-agent scheduling and planning activities that involve drilling and excavation.

Several observations about the BGM system should be noted. First, although the boundary segmentation process has been designed to extract contours from sparse spatial patterns, it can take as input any sensible closed contour generated by humans or computers, such as lithological boundary detection




FIGURE 15. Extracted surfaces for geozone g_1 (orange) and g_2 (blue) conform to the local geological structure where one is stacked on top of the other. Two vantage points from above and below in the top and bottom panels.

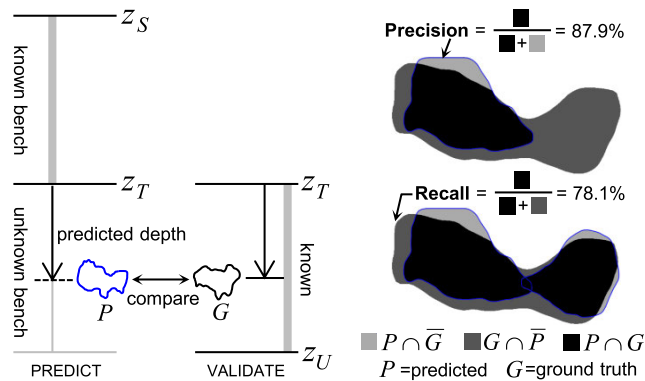


FIGURE 16. Concept of precision and recall for spatial regions.

using rock face image segmentation [53] and 2D decision boundaries from Gaussian process implicit surfaces [54] or SVM. Second, for spatial correspondence, it is possible to have humans in the loop to provide interactive estimates and perform component alignment in difficult cases. Third, the contour metamorphosis step may be preceded by geometric operations such as affine transform. For instance, a rotation between the source and target components can be estimated in the Fourier domain using the technique described by Nagashima [28]. Fourth, the directional estimates generated by the model can feedback into the system and act as




FIGURE 17. Precision and recall rates for prediction: (gray) proposed model, (black) scenario where slope information is unavailable.

a prior during region association. Finally, when data becomes more abundant, machine learning techniques can potentially be used to estimate certain parameters relating to surface dynamics (e.g., the anticipated lateral movement for certain regions).

VII. CONCLUSION

This paper presented a framework for modeling geological boundaries using differential geometry. The objective was to create subsurfaces from sparse spatial patterns and obtain coherent directional predictions along the boundary of the extracted surfaces. Under the model, the precision and recall rate for contour prediction improved on average by 15–20%. For boundary extraction, an edge map synthesis procedure was described. This converted sparse non-uniform data points to an image representation and facilitated the use of PDE-based techniques which operate on dense uniform 2D arrays. Gradient vector field and active contours were used to obtain regularized contours from unordered edge pixels. For spatial correspondence, region association and component alignment problems were examined, translation estimates were obtained using FFT cross-correlation under spatial constraints. Multi-source multi-target component relationships were explored using intersection graphs;

TABLE 4. List of appendixes (see supplemental material).

A	Theoretical treatment of active contours & gradient vector field
B	Foundations for contour metamorphosis

TABLE 5. List of algorithms (see supplemental material).

	Description	Ref. section
1	Active contour evolution in GVF	Appendix A
2	Edge map synthesis	§ II-D
3	Signed distance function	§ III-B1
4	Single-source single-target correspondence	§ III-B1
5	Port allocation to handle spatial contention	§ III-B2
6	Multi-source single-target correspondence	§ III-B2
7	Single-source multi-target correspondence	§ III-B3
8	Multi-source multi-target correspondence	§ III-B4
9	Unified spatial correspondence strategy	§ III-B4
10	Wavefront curvelets backtracking	§ IV-C, Fig. 12
11	Wavefront curvelets branching analysis	§ IV-C, Fig. 13
12	Assembling particle trajectories (merging existing and new trajectories from curvelets)	§ IV-C, Fig. 13

strategies for obstacle avoidance were formulated from a resource allocation perspective. For contour metamorphosis, local surface deformation was modeled using PDE (described in Appendix B). Surface slopes were estimated using normalized particle trajectories. The branching phenomenon was described: branching occurs during contour morphing when there is a significant mismatch in curvature between the source and target boundary segments. A curvelet backtracking algorithm was proposed to recover information lost during particle advection and thus overcome tracking failure caused by branching. In essence, the overall solution dealt with sparsity (irregularity), spatial constraints, and branching-induced tracking failure.

In a *Nature* editorial, an opinion piece [55] characterized interdisciplinary research as a synthesis of different approaches into something unique. This captures the essence of this study which is to re-imagine how concepts from established areas can be synthesized and applied to a different field. This study had a narrow focus on estimating and exploiting the directionality of subsurface boundaries for more targeted drilling and exploration. Its uniqueness stems from applying differential geometry to sparse data. A different blend of technologies may be appropriate in a different setting where automation involves something different. Whilst research papers in engineering and the natural sciences have increasingly cited work outside their own disciplines [56], some areas remain unexplored. These areas can benefit from fresh perspectives and offer opportunities for meaningful collaboration [57]. It is hoped this work can encourage more people to engage in and apply their expertise to various emerging and non-traditional fields.

APPENDIXES

Appendixes are included as supplemental material. The subject matters are outlined in Table 4.

ALGORITHMS

Algorithms are included as supplemental material. Table 5 lists the algorithms and referenced sections.


ACKNOWLEDGMENT

Prof. S. Sukkarieh and the reviewers are thanked for their suggestions and considered opinion which improved the quality of this manuscript.

REFERENCES

- [1] S. J. Buckley, J. A. Howell, H. D. Enge, and T. H. Kurz, "Terrestrial laser scanning in geology: Data acquisition, processing and accuracy considerations," *J. Geol. Soc.*, vol. 165, no. 3, pp. 625–638, 2008.
- [2] T. Frank, A.-L. Tertois, and J.-L. Mallet, "3D-reconstruction of complex geological interfaces from irregularly distributed and noisy point data," *Comput. Geosci.*, vol. 33, no. 7, pp. 932–943, Jul. 2007.
- [3] O. Nilsson, D. Breen, and K. Museth, "Surface reconstruction via contour metamorphosis: An Eulerian approach with Lagrangian particle tracking," in *Proc. IEEE Visualization*, Oct. 2005, pp. 407–414.
- [4] D. Adalsteinsson and J. A. Sethian, "A fast level set method for propagating interfaces," *J. Comput. Phys.*, vol. 118, no. 2, pp. 269–277, 1995.
- [5] K. B. Sprague and E. A. De Kemp, "Interpretive tools for 3-D structural geological modelling part II: Surface design from sparse spatial data," *Geoinformatica*, vol. 9, no. 1, pp. 5–32, 2005.
- [6] H. Mitášová and J. Hofierka, "Interpolation by regularized spline with tension: II. Application to terrain modeling and surface geometry analysis," *Math. Geol.*, vol. 25, no. 6, pp. 657–669, 1993.
- [7] O. Kaufmann and T. Martin, "3D geological modelling from boreholes, cross-sections and geological maps, application over former natural gas storages in coal mines," *Comput. Geosci.*, vol. 34, no. 3, pp. 278–290, 2008.
- [8] J.-L. Mallet, "Discrete smooth interpolation in geometric modelling," *Comput.-Aided Des.*, vol. 24, no. 4, pp. 178–191, 1992.
- [9] A. Zanchi, S. Francesca, Z. Stefano, S. Simone, and G. Graziano, "3D reconstruction of complex geological bodies: Examples from the Alps," *Comput. Geosci.*, vol. 35, no. 1, pp. 49–69, 2009.
- [10] G. Caumon, P. Collon-Drouaillet, C. Le Carlier de Veslud, S. Viseur, and J. Sausse, "Surface-based 3D modeling of geological structures," *Math. Geosci.*, vol. 41, no. 8, pp. 927–945, 2009.
- [11] B. De Waele, M. Lacorde, M. Cunningham, and B. Jupp, "Understanding geology and structure: An essential part of mineral resource estimation," *ASEG Extended Abstr.*, vol. 2018, no. 1, pp. 1–8, 2018.
- [12] J. Dirstein, T. Rudge, R. Li, and A. Stanley, "Insights from the automated extraction of surfaces from the Bunda 3D seismic survey," in *Proc. West Austral. Basins Symp.*, Aug. 2013, pp. 18–21.
- [13] N. Voorhies, S. Bowen, T. Battenhouse, R. Porges, and T. Fox, "From 2D cross-sections to 3D model: A toolset for integrated data management, modelling, and visualization," in *Proc. Three-Dimensional Geol. Mapping, Workshop Extended Abstr.*, Geol. Soc. Amer. Annu. Meeting, R. C. Berg, H. A. Russell, and L. H. Thorleifson, Eds. Portland, Oregon: Illinois State Geological Survey, 2009. [Online]. Available: <https://library.isgs.illinois.edu/Pubs/pdfs/ofs/2009/ofs2009-04.pdf>
- [14] R. C. Berg, S. J. Mathers, H. Kessler, and D. A. Keefer, "Synopsis of current three-dimensional geological mapping and modeling in Geological Survey organizations," *Illinois State Geol. Surv. Circular*, vol. 578, p. 104, 2011. [Online]. Available: <https://nora.nerc.ac.uk/17095/1/c578.pdf>
- [15] M. D. Lindsay, M. W. Jessell, L. Ailleres, S. Perrouty, E. de Kemp, and P. G. Betts, "Geodiversity: Exploration of 3D geological model space," *Tectonophysics*, vol. 594, pp. 27–37, May 2013.
- [16] B. Lin, L. Zhou, G. Lv, and A.-X. Zhu, "3D geological modelling based on 2D geological map," *Ann. GIS*, vol. 23, no. 2, pp. 117–129, 2017.
- [17] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1988.
- [18] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. J. Comput. Vis.*, vol. 22, no. 1, pp. 61–79, 1997.
- [19] J. Ivins and J. Porrill, "Everything you always wanted to know about snakes (but were afraid to ask)," *Artif. Intell. Vis. Res. Unit, Univ. Sheffield, U.K., AIVRU Tech. Memo* 86, Mar. 2000.
- [20] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 359–369, Mar. 1998.

- [21] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, nos. 1–3, pp. 185–203, Aug. 1981.
- [22] D. E. Breen and R. T. Whitaker, "A level-set approach for the metamorphosis of solid models," *IEEE Trans. Vis. Comput. Graphics*, vol. 7, no. 2, pp. 173–192, Apr. 2001.
- [23] K. Museth, D. E. Breen, R. T. Whitaker, S. Mauch, and D. Johnson, "Algorithms for interactive Editing of level set models," *Comput. Graph. Forum*, vol. 24, no. 4, pp. 821–841, 2005.
- [24] M. B. Nielsen and K. Museth, "Dynamic tubular grid: An efficient data structure and algorithms for high resolution level sets," *J. Sci. Comput.*, vol. 26, no. 3, pp. 261–299, 2006.
- [25] M. Bertalmio, G. Sapiro, and G. Randall, "Morphing active contours," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 7, pp. 733–737, Jul. 2000.
- [26] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. Comput. Phys.*, vol. 79, no. 1, pp. 12–49, 1988.
- [27] D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang, "A PDE-based fast local level set method," *J. Comput. Phys.*, vol. 155, no. 2, pp. 410–438, 1999.
- [28] S. Nagashima, K. Ito, T. Aoki, H. Ishii, and K. Kobayashi, "A high-accuracy rotation estimation algorithm based on 1D phase-only correlation," in *Image Analysis and Recognition*. Berlin, Germany: Springer, 2007, pp. 210–221.
- [29] M. Balamurali and A. Melkumyan, "T-SNE based visualisation and clustering of geological domain," in *Proc. Int. Conf. Neural Inf. Process.* Kyoto, Japan: Springer, 2016, pp. 565–572.
- [30] R. Ardon, L. D. Cohen, and A. Yezzi, "Fast surface segmentation guided by user input using implicit extension of minimal paths," *J. Math. Imag. Vis.*, vol. 25, no. 3, pp. 289–305, Oct. 2006.
- [31] R. Ardon, L. D. Cohen, and A. Yezzi, "A new implicit method for surface segmentation by minimal paths in 3D images," *Appl. Math. Optim.*, vol. 55, no. 2, pp. 127–144, 2007.
- [32] R. Goodfellow, F. A. Consuegra, R. Dimitrakopoulos, and T. Lloyd, "Quantifying multi-element and volumetric uncertainty, coleman McCreeley deposit, Ontario, Canada," *Comput. Geosci.*, vol. 42, pp. 71–78, May 2012.
- [33] P. Tahmasebi and A. Hezarkhani, "A hybrid neural networks-fuzzy logic-genetic algorithm for grade estimation," *Comput. Geosci.*, vol. 42, pp. 18–27, May 2012.
- [34] M. J. Cracknell and A. M. Reading, "Geological mapping using remote sensing data: A comparison of five machine learning algorithms, their response to variations in the spatial distribution of training data and the use of explicit spatial information," *Comput. Geosci.*, vol. 63, pp. 22–33, Feb. 2014.
- [35] K. L. Silversides and A. Melkumyan, "Gaussian processes based fusion of multiple data sources for automatic identification of geological boundaries in mining," in *Proc. Int. Conf. Neural Inf. Process.* Kyoto, Japan: Springer, 2016, pp. 294–301.
- [36] A. Karpatne, I. Ebert-Uphoff, S. Ravela, H. A. Babaie, and V. Kumar, "Machine learning for the geosciences: Challenges and opportunities," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 8, pp. 1544–1554, Aug. 2018.
- [37] B. Jafarsteh, N. Fathianpour, and A. Suárez, "Comparison of machine learning methods for copper ore grade estimation," *Comput. Geosci.*, vol. 22, no. 5, pp. 1371–1388, 2018.
- [38] L. Shapiro and G. Stockman, *Computer Vision*. Upper Saddle River, NJ, USA: Prentice-Hall, 2002, ch. 3.4.
- [39] K. Suzuki, I. Horiba, and N. Sugie, "Linear-time connected-component labeling based on sequential local operations," *Comput. Vis. Image Understand.*, vol. 89, no. 1, pp. 1–23, Jan. 2003.
- [40] D. Koller, J. Weber, and J. Malik, "Robust multiple car tracking with occlusion reasoning," in *Computer Vision—ECCV*. 1994, pp. 189–196.
- [41] M. Isard and A. Blake, "Condensation—Conditional density propagation for visual tracking," *Int. J. Comput. Vis.*, vol. 29, no. 1, pp. 5–28, 1998.
- [42] C. Rasmussen and G. D. Hager, "Probabilistic data association methods for tracking complex visual objects," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 560–576, Jun. 2001.
- [43] G. Tsechpenakis, K. Rapantzikos, N. Tsapatsoulis, and S. Kollias, "A snake model for object tracking in natural sequences," *Signal Process. Image Commun.*, vol. 19, no. 3, pp. 219–238, 2004.
- [44] O. Williams, A. Blake, and R. Cipolla, "Sparse Bayesian learning for efficient visual tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1292–1304, Aug. 2005.
- [45] B. Stenger, A. Thayananthan, P. H. S. Torr, and R. Cipolla, "Model-based hand tracking using a hierarchical Bayesian filter," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 9, pp. 1372–1384, Sep. 2006.
- [46] Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi, "Tracking deforming objects using particle filtering for geometric active contours," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 8, pp. 1470–1475, Aug. 2007.
- [47] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2002.
- [48] P. F. Felzenszwalb and D. P. Huttenlocher, "Distance transforms of sampled functions," *Theory Comput.*, vol. 8, no. 1, pp. 415–428, 2012.
- [49] J. Tomek, R. A. Burton, and G. Bub, "Ccoffinn: Automated wave tracking in cultured cardiac Monolayers," *Biophys. J.*, vol. 111, pp. 1595–1599, Oct. 2016.
- [50] R. Singh, I. Pavlidis, and N. P. Papanikolopoulos, "Recognition of 2D shapes through contour metamorphosis," in *Proc. Int. Conf. Robot. Automat.*, Apr. 1997, pp. 1651–1656.
- [51] C. Zhang and T. Chen, "Efficient feature extraction for 2D/3D objects in mesh representation," in *Proc. Int. Conf. Image Process.*, vol. 3, Oct. 2001, pp. 935–938.
- [52] N. Ahsan, S. Scheding, S. T. Monteiro, R. Leung, C. McHugh, and D. Robinson, "Adaptive sampling applied to blast-hole drilling in surface mining," *Int. J. Rock Mech. Mining Sci.*, vol. 75, pp. 244–255, Apr. 2015.
- [53] Y. Vasuki, E.-J. Holden, P. Kovese, and S. Micklethwaite, "An interactive image segmentation method for lithological boundary detection: A rapid mapping tool for geologists," *Comput. Geosci.*, vol. 100, pp. 27–40, Mar. 2017.
- [54] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," in *Gaussian Proc. Pract.*, 2007, pp. 1–4.
- [55] "Mind meld, in interdisciplinary," *Nature Editorial*, vol. 525, no. 7569, pp. 289–290, 2015, doi: [10.1038/525289b](https://doi.org/10.1038/525289b).
- [56] R. Van Noorden, "Interdisciplinary research by the numbers," *Nature*, vol. 525, no. 7569, pp. 306–307, 2015.
- [57] R. R. Brown, A. Deletic, and T. H. Wong, "Interdisciplinarity: How to catalyse collaboration," *Nature*, vol. 525, no. 7569, pp. 315–318, Sep. 2015, doi: [10.1038/525315a](https://doi.org/10.1038/525315a).



algorithm development. His doctoral studies were supported by the APA Australian Government Scholarship.

• • •