

Batch automated image processing of 2D seismic data for salt discrimination and basin-wide mapping

Scott Morris¹, Shuang Li², Tony Dupont¹, and John D. Grace¹

ABSTRACT

We have explored the technical utility of analyzing massive sets of digital 2D seismic data, collected and processed in dozens of different surveys, conducted more than 25 years ago, using batch, automated and unsupervised pattern recognition techniques to produce a basin-wide map of the top of salt. This workflow was developed for the United States portion of the Gulf of Mexico to detect top-salt boundaries on 2D poststack migrated lines. Texture-based attributes as well as novel, reflector-based attributes were used to discriminate between salt and nonsalt on each seismic line. Explicit measures of accuracy were not calculated because the data are unlabeled, but an assessment of confidence was used to score the boundaries. The depth to the top of the salt was estimated for more than 67% of the study area (278,000 km² or 107,000 mi²), 17% of the study area had insufficient data for processing and analysis, and 16% of the area did not meet confidence requirements for inclusion. The final results compared well with published maps of salt and the locations of salt-trapped fields. Reliable mapping of salt deeper than 6 s two-way time could not be achieved with this data set and approach because many seismic images had indistinguishable features at this depth. The computing time was greater than linear in the number of lines, but parallelization and changes in hardware configuration could reduce the run time of about three weeks to about three days.

INTRODUCTION

Seismic technology's advancement over the past three decades, particularly the transition from 2D to 3D, has seemingly rendered

enormous volumes of vintage 2D data obsolete. We offer two hypotheses to challenge that conclusion. First, applying an innovative combination of image processing techniques, in batches, produces valuable new information from these legacy assets at very low marginal costs. Second, changing the scale of intensive analysis from the prospect/field level to a whole region, or a basin, further leverages the return to investment in automated interpretation.

To test these propositions, we map basin-wide top of salt across the U.S. portion of the Gulf of Mexico (GOM). We applied a workflow of image processing, analysis of reflectors, and macroediting techniques to a set of 8043 2D poststack, migrated SEG-Y lines that cover approximately 278,000 km² (107,000 mi²) of the northern shelf and slope. They were shot by about a dozen operators and contractors in 82 surveys, *all more than 25 years ago*, between 1981 and 1992. This trove of disparate data therefore reflects the character of many of the seismic legacy assets that companies hold, as well as some of the seismic data now being released into the public domain (e.g., Common Data Access [CDA, 2018] in the United Kingdom and the United States Geological Survey [Triezenberg et al., 2016] in the United States).

Our central scientific enterprise, constructing a basin-wide product using thousands of lines of seismic data, accommodating disparities in their age and provenance — while minimizing cost — required renovating previous approaches and departing from some analytic traditions. The geophysical literature on automated salt discrimination is dominated by examples developed on relatively small, field data sets, often with lines on which salt was known to exist, *a priori*. (Alternatively, some research has been conducted on larger data sets; however, these were synthetic data sets, which are also useful as proof-of-concept tests but do not face the problems of application to field seismic data. See Araya-Polo et al. (2017) for methods developed for fault identification, tested on synthetic data, using an approach that could be applied to salt.) These small-sample, proof-of-concept experiments have been highly valuable. After more than 20 years of research, they founded a new paradigm for geophysical

Manuscript received by the Editor 30 July 2018; revised manuscript received 17 July 2019; published ahead of production 4 September 2019; published online 30 October 2019.

¹Earth Science Associates, Long Beach, California, USA. E-mail: scott@earthsci.com; tony@earthsci.com; john@earthsci.com.

²University of Southern California, Earth Science Associates and Graduate Program in Applied Mathematics, Los Angeles, California, USA. E-mail: lishuang@usc.edu; tony@earthsci.com.

© 2019 Society of Exploration Geophysicists. All rights reserved.

analysis. However, rather than locating salt on a line where it had already been found manually, we investigate the scalability and robustness of these methods and a workflow for a basin full of seismic.

Finally, the scale of our approach required special attention to two issues that do not arise with small test sets. Naturally, massive data sets always require consideration of computational complexity; a regional or basin-wide process must run within a practical length of time. Moreover, of great practical importance, analyzing thousands of lines meant hand-checking accuracy by line was impossible. Therefore, methods for automatic evaluation of confidence in the estimated boundaries were developed to convey certainty in the analysis to users.

We address four basic questions. Is our methodology, and its component algorithms, technically effective for regional to basin-wide mapping of salt? Central to that question is the adequacy of our combined analyses of seismic texture and reflector geometry for salt/nonsalt classification. To the basic requirement that the method works, we add: Can this approach be generalized — scaled to batch analysis of thousands of lines — and thereby be usefully applied to a region or basin? Is this methodology robust in the face of the problems of old data, collected from a large number of disparate surveys? Finally, are the marginal costs to implement our approach low enough to motivate exploiting the countless, enormous legacy seismic assets, which once cost billions of dollars to acquire, but are now otherwise *free* but unused?

METHODOLOGY

Two classes of attributes were used to discriminate between salt and nonsalt in the seismic images. First, seismic texture, the characteristic patterns of reflection amplitudes defined by the magnitude and variation of neighboring samples at a given spatial location (Gao, 2003), was exploited. Texture is the primary attribute in many image segmentation problems (Bhosle and Pawar, 2013) and has been used frequently to detect salt bodies (Hegazy and AlRegib, 2014; Shafiq et al., 2017b). Second, reflector analysis used the well-known relationship between the geometry and distribution of reflectors and salt domes. Estimated dip attributes (Halpert and Clapp, 2008) have been used to exploit this relationship (Berthelot et al., 2013), but this method extracts whole reflectors of a seismic image and categorizes them rather than relying on texture analysis.

The texture-based attributes were calculated using the gray-level co-occurrence matrix (GLCM) methodology, originally developed by Haralick et al. (1973), to help partition areas of salt and nonsalt. The development of GLCM in geophysics, and its fairly wide application in discrimination, are reviewed by Gao (2011) and Berthelot et al. (2013). We focused on four salt-discriminatory GLCM statistics: contrast, homogeneity, entropy, and dissimilarity.

The texture-based discrimination was supplemented by reflector analysis. Grouping contiguous pixels of similar intensity in the seismic image allows reflectors to be treated as individual logical objects, rather than many pixels. The reflectors, easily subset on length and angle, facilitate the search for specific geometric properties observed around salt domes. For one measure, the reflectors were used to directly estimate the salt/nonsalt boundary. In the second measure, reflectors were segregated by geometric characteristics associated with locally bedded (nonsalt) rock, and their densities mapped, which provided evidence of where salt is not likely to be present.

The outputs of each of these methods were combined into a *salt score* for each pixel in a line, which was then subject to a two-group clustering algorithm (Otsu, 1979). The resulting raw black-and-white

image of the line was then edited with morphological tools to (1) filter out the smallest estimated salt bodies (presumptive false positives), (2) fill the smallest holes in estimated salt bodies (presumptive false negatives), and (3) smooth boundaries to more closely approximate geologically acceptable interpretations (Wallet and Pepper, 2013).

The focus then shifted to evaluation of confidence in the estimated salt/nonsalt boundaries measuring the change in texture (CIT) along the boundary. This method takes inspiration from the gradient of texture (GOT), which is commonly used in image analysis to evaluate boundaries between textures, as described by Shafiq et al. (2017a) and Hegazy et al. (2015). The average CIT of each salt polygon became a grade of confidence in the estimated boundary; this influenced the decision to retain or reject each candidate boundary. The average CIT scores for each boundary were also reflected as points along the line's navigation data coincident with the lateral extent of each estimated salt boundary to produce regional maps of confidence.

In the final step, each line and its estimated salt boundaries were converted from time to depth using empirically estimated velocity functions. Results in depth are expected by geologists and engineers. The final products were confidence-marked top of salt maps in time and depth and 3D surfaces for use in 3D geographic information systems. The project workflow is outlined in Figure 1.

Data

Data for our experiment were drawn from publicly available seismic surveys submitted to the U.S. government by operators in the U.S. federal portion of the GOM (United States Bureau of Ocean Energy Management [USBOEM], 2018). The 8043 2D time sections totaled approximately 450,000 km (280,000 mi) in length, averaged 56 km (35 mi) long, and typically extended to approximately 8 s in two-way time. The U.S. government holds these records confidential for 25 years before public release. All surveys used here were shot between 1981 and 1992. They are, by any measure, old data.

We illustrate our methodology using part of one line, “w-86-314,” from survey “1986 namss.B-39-86-LA.msc.airgun,” a 51.4 km (31.9 mi) line, with a maximum recorded time of 7.5 s (Figure 2a). It was shot in 1986 by Western Geophysical from northwest to southeast across the southwestern part of the Ship Shoal protraction area of the GOM (Figure 2b). It covers the Ship Shoal 154 field, trapped on a salt dome, with a latest estimated ultimate recovery of 132 million barrels of oil equivalent (USBOEM, 2018).

Texture analysis

The amplitude records in each SEG-Y file were read and mapped to a bin based on the number of standard deviations ($|\sigma(A)|$) away from a line-wise mean value, which, in almost all cases, was zero. We binned them into a seven-level gray scale (Table 1). These boundaries reflect a nonlinear transformation in which very low amplitude values (which can produce false-positive salt classifications) were given a more detailed breakdown in gray levels to help limit these errors. Conversely, very high absolute values of amplitude were compressed because these provide diminishing useful information for salt discrimination.

The seismic texture of salt (particularly in domes) is often obvious on visual inspection: Reflectors in salt bodies are typically and distinguishably weak, short, chaotically oriented, and sparse compared to surrounding stratified rock. Our analysis attempts to approximate the visual differentiation that comes naturally to an expert interpreter.

In this, we extend previous work using GLCM analysis of seismic texture as discussed in Wu (2016), Amin et al. (2015), Wang et al. (2015), Hegazy and Alregib (2014), and Berthelot et al. (2013). Our implementation used the *GLCM* library of *R* (R Core Team, 2013), documented by Zvoleff (2016). It consists of two steps executed for each pixel in each seismic section: Build the GLCMs for all pixels, and compute statistics for each pixel's GLCM. For greater detail on implementing GLCM, see Appendix A.

The GLCM statistics become attributes of each pixel and are therefore mappable for visual assessment of classification accuracy. In experiments over subsets of lines from different surveys, vintages, and regions, four GLCM statistics were chosen to maximize the accuracy of classifying salt versus nonsalt areas and minimize problems of false positives: contrast α^1 , homogeneity α^2 , dissimilarity α^3 , and entropy α^4 Figure 3a–3d).

The functions for three of the four GLCM statistics chosen are closely related. Dissimilarity and contrast vary only by the weighting factor for the distance between each probability cell (p_{ij}) in the GLCM and its principal diagonal. For dissimilarity, the weight is linear, and for contrast, it is squared. Homogeneity and contrast are inversely correlated but not exactly. Both are also separate functions of GLCM energy, which varies with the number and the likeness of nonzero elements in the GLCM (Baraldi and Parmiggiani, 1995). Homogeneity is not an exact inverse of dissimilarity for the same reason.

The values of each statistic were normalized to a [0,1] scale. Additionally, all of the GLCM statistics, with the exception of homogeneity, were converted from their original equations by defining them as $(1 - \alpha_{mn}^k)$. This makes all the statistics comparable with 1 associated with the presence of salt and 0 the absence of salt.

Reflector analysis

To estimate the two reflector-based components of the analysis, the data from the SEG-Y were divided into two sets: positive and negative amplitudes. Then, for each of those two sets, the observations were binned by standard deviation using the classes in Table 1. All observations in the positive and negative amplitude sets greater than or equal to a threshold of 0.25σ were retained and labeled lower intensity reflectors. This process was repeated, using a cut-off $\geq 1\sigma$, to create a set that was labeled higher-intensity reflectors. Finally, we applied the *Raster to Polygon* tool in ArcGIS-Pro (ESRI, 2018) to both sets to group the contiguous pixels meeting the intensity thresholds so reflectors are treated as single logical objects.




Figure 1. The general workflow for the application of batch discrimination of salt in 2D seismic and the creation of the top of salt maps in time and depth.




Figure 2. (a) Grayscale image of a portion of the time section of the 'w-86-314' line in survey '1986 namss.B-39-86-LA.msc.airgun' running northwest to southeast, left to right. (b) Mapping of the line and fields in the area (light gray) and the SS154 field (dark gray). The numbered squares are lease boundaries. Ship Shoal (SS).

The first reflector-based discriminator exploits an often-observed regularity in diapiric salt on 2D seismic: Salt bodies are typically bounded on both sides by relatively short, high-dip, usually intense reflectors. When the records in the salt body itself are eliminated, these bounding reflectors form chevrons going up the salt dome vertically.

In creating this parameter, the set of higher-intensity reflectors was subset on two criteria: (1) those with the absolute value of the dip $>40^\circ$ or (2) those reflectors from the higher intensity set that were in the longest 2.5% of the distribution of reflector lengths. The latter group was included because, on visual examination, we observed that sometimes on the boundaries of domes, very long reflectors actually extend from one side of the dome, over the top and down the other side, helping to define the boundary. The result of this subsetting is shown in Figure 4a; the salt dome is very evident.

An algorithm, written in native *R*, was created to find the area bounded by these reflectors. It begins by reading row-wise from left to the right, searching for a positively sloped reflector. Once found, the search continues to the right for a *matching* negatively sloped or very strong reflector. This algorithm was then applied on the same row in the opposite direction. The pixels in between these markers (i.e., the bounding chevrons) were turned to 1, and those

outside the chevrons on that line turned to 0. The analysis was repeated on the next row down until the entire line was processed. The result of this processing is shown in Figure 4b.

After the raw image was created, it was edited with morphological tools (see below), eliminating very small features and smoothing edges (Figure 4c). Every pixel was labeled 1, indicating salt, or 0 for its absence, and constituted the fifth salt measure α_{mn}^5 , as shown in Figure 4c.

The final salt measure, α_{mn}^6 , is created from the set of long, low-dip reflectors subset by removing the shortest 10th percentile of reflectors. The density of the remaining reflectors is calculated in *R* based on a 51×51 moving window and is scaled [0:1] with 1 indicating no reflectors in the window (Figure 4d). Particularly, for diapiric salt on the shelf, the long, low-dip reflectors of bedded rock surrounding domes characterize where salt is *not*. Therefore, α_{mn}^6 near 1 indicates the presence of salt.

Salt score and partitioning

Completing the texture and reflector analyses for a single seismic line produced six measures on salt for each pixel. Five of them (the four GLCM statistics [$\alpha_{mn}^1, \dots, \alpha_{mn}^4$] and the long-reflector density raster [α_{mn}^6]) were continuous variables from zero to one. The other reflector measure, being inside or outside matching high-dip reflectors, was treated as a two-state (0/1) categorical variable α_{mn}^5 .

For each pixel in the line, a *salt score* (s_{mn}) was computed as the arithmetic average of the six salt measures ($\alpha_{mn}^1, \dots, \alpha_{mn}^6$). The raw salt score for the study line is shown in Figure 5a. To partition the salt score raster into areas of salt and nonsalt, we applied a clustering algorithm (Otsu, 1979). It reflects a common approach to statistical discriminant analysis: assign samples (pixels) to one of two classes to maximize the between-class variance and minimize the within-class variance. This can be algebraically reduced to maximize the ratio of the between-class variance to the total variance of pixels' grayscale values over the entire image. The raw partitioning of the study section into salt/nonsalt is shown in Figure 5b.

Morphological processing

The raw salt/nonsalt partition in Figure 5b is an image in need of editing to eliminate the effects of processing. This included removing very small estimated salt bodies and cleaning the boundaries of the larger estimated salt bodies, principally to remove high-frequency noise. The opening and closing tools were used for these tasks, as described by Shafiq et al. (2015), Wallet and Pepper (2013), and Urbach and Wilkinson (2008). The result is shown in Figure 6a. Then, small holes of *nonsalt* were filled in within bodies of large salt features. All three methods were implemented in the *EBImage* package in *R* (Pau et al., 2010). After all morphological processing was complete, the final time boundaries were extracted (Figure 6b).

Table 1. Transformation from SEG-Y amplitude values to grayscale levels.

Gray level	Bin	Gray value	Gray level	Bin	Gray value
1	$\sigma(A) \leq 0.1$	White	5	$1 > \sigma(A) \leq 1.5$	80% gray
2	$0.1 > \sigma(A) \leq 0.25$	15% gray	6	$1.5 > \sigma(A) \leq 2$	90% gray
3	$0.25 > \sigma(A) \leq 0.5$	30% gray	7	$\sigma(A) > 2.0$	Black
4	$0.5 > \sigma(A) \leq 1$	60% gray			

Note: $|A|$, absolute value of amplitude; σ , standard deviation.




Figure 3. The GLCM statistics on a subset of Figure 2a around the salt dome: (a) contrast, (b) homogeneity, (c) dissimilarity, and (d) entropy. Each figure has the scale 0–1 (black to white), with 0 interpreted as nonsalt and 1 as salt.

Macroediting

When the analysis, through morphological processing, was complete, the results in time were mapped for a Gulf-wide top of salt map. Inspection revealed some problems at a macrolevel (i.e., at the survey or regional levels). These fell into four groups:

- 1) Local outliers: Certain locations contain multiple seismic lines (often from disparate surveys) in very close proximity. In some of these areas, the results from one line may be different from the others, which create local outliers. A cluster analysis was performed in ArcGIS Pro using the *Anselin Local Moran's I* tool (ESRI, 2018) to locate these outliers and remove them.
- 2) Survey-level acquisition/processing problems: In some surveys, all lines shared characteristics that interfered with our process. A few surveys were visually *too dark* relative to the others, apparently due to different handling of the range of amplitudes. The change in some of the components of the GLCM analysis cured, or at least alleviated, some of these problems.
- 3) Local geologic anomalies within surveys: On inspection, some of the apparent salt bodies seemed more likely to be expressions of other geologic features. This occurred around the mouth of the current course of the Mississippi River; it was also mirrored at two previous Mississippi courses as they came offshore. They may be distributary mouth bars or related deltaic facies that produced, from a salt discrimination viewpoint, false positives. The questionable data within the lines affected were removed.
- 4) Survey-level problems of unknown provenance: In a few cases, large areas of contiguous shallow *salt* were mapped, apparently correlated with survey boundaries. Inspection yielded no corresponding salt, and the source of the irregularities could not be identified. No surgical excision or algorithm modification could suitably resolve the problem, so the surveys were dropped (e.g., from the Brazos protraction area on the Texas shelf to the south).

Additionally, many false positives were created by salt boundaries that were far too small to be anything more than processing artifacts not removed by the previous morphological operations. The smallest 75% of the boundary lines were removed. This strict rule certainly removed some valid boundaries; however, the cost of removing some good boundaries versus the benefit to also remove many bad boundaries is a good trade-off when creating a basin-wide top of salt map.




Figure 4. The intermediate and final results of analysis of the reflector data. (a) Mapping of the strong reflectors in the study line. High-dip reflectors with positive slopes (up to the right) are shown in white; high-dip lines with negative slopes are shown in light gray. Strong low-dip lines are shown in dark gray. (b) The assignment of the space between paired high-dip reflectors to white. (c) The result of morphological editing of the raw image (b). (d) The density of long, and generally, low-dip reflectors, indicating the absence of salt.




Figure 5. (a) The raw salt score computed from the six classification measurements ($\alpha^1, \dots, \alpha^6$). (b) The transformation of the raw salt score image in (a) into a two-state clustering through the application of Otsu's (1979) algorithm.




Figure 6. (a) The effect of the application of the opening and closing morphological editing tools (the *fill* tool is not separately shown because its role in this line was to fill the two small holes in the interior of the salt in [a]). (b) The final, postmorphological editing, salt boundary on the time image of the study line.

Measuring confidence

In processing more than 8000 seismic lines, with no pretense to hand check the quality and reliability of each one individually, communicating certainty in the analysis was critical for building a technologically useful process. Therefore, an automated assessment was applied to every salt boundary estimated in the study. At the next level of aggregation, the average certainty score for each boundary in a line was projected to the points overlying its lateral extent onto the navigation data of the line for regional mapping.

We adopted two screening criteria and applied them to each estimated salt boundary. The principal measure was a simple CIT, inspired by the GOT method. The GOT is used frequently in detecting salt bodies using texture attributes (Hegazy et al., 2015) with a wide variety of texture dissimilarity functions (Shafiq et al., 2017b). The CIT measures the change in pixel intensity at a given pixel by using a vertical and horizontal search window. Although the CIT can be computed for any pixel (with an exception near edges), here it is only calculated along the estimated salt boundary. The CIT measure gives equal weight to all pixels in the window by performing a convolution using a large (11×51 or 51×11) filter mask. Mathematically, it is defined as follows:

$$C_x = \begin{bmatrix} 1 & \dots & 1 & 0 & -1 & \dots & -1 \\ 1 & \dots & 1 & \dots & -1 & \dots & -1 \\ 1 & \dots & 1 & 0 & -1 & \dots & -1 \end{bmatrix} * \mathbf{W} \quad (1)$$

11×51

$$C_y = \begin{bmatrix} 1 & \dots & 1 & 0 & -1 & \dots & -1 \\ 1 & \dots & 1 & \dots & -1 & \dots & -1 \\ 1 & \dots & 1 & 0 & -1 & \dots & -1 \end{bmatrix}^T * \mathbf{W}, \quad (2)$$

51×11

where C_x is the change in the x -direction, C_y is the change in the y -direction, $\mathbf{W}$ is the appropriately sized window of pixel intensity values around the target pixel, and $*$ is the operator for the convolution between two matrices. We then define the CIT $C = \max(|C_x|, |C_y|)$ rather than using the usual L2 norm because a well-placed boundary will typically have very different changes in the x - and y -directions




Figure 7. The principle behind the CIT measure. In (a), a bisected rectangle is centered on a pixel in the estimated boundary of a salt dome. The intensity of pixels within the rectangle outside the boundary is clearly different from that of the pixels interior to the boundary. In (b), the average pixel intensity of both sides of the rectangle is nearly the same, so the boundary is removed.

using this method. A large filter mask was used because the CIT is not used to determine the salt boundary (only to measure the confidence in the one estimated) and allows an imperfectly placed (though still reasonable) salt boundary to receive a good CIT score and be retained in the analysis. This specific size was chosen so the width/length of the window matched the width/length used in the GLCM analysis. The spatial coarseness of this method reinforces the caveat that our analysis is regional and not intended for high-resolution placement of individual salt boundaries.

Conceptually, a small rectangle is centered on a pixel in a boundary, and the rectangle is divided in half at the boundary (Figure 7a and 7b). The sum of pixel intensities is computed in each half of the rectangle (i.e., inside and outside the boundary) and the difference taken between the two. This is repeated using horizontal and vertical rectangles, and the rectangle with the larger difference is used. A significant difference in the pixel intensity between the two halves of the rectangle evidences a significant difference in seismic textures across the boundary and is taken as evidence for an appropriately located boundary.

The CIT is calculated for each pixel along the boundary, and the average CIT is assigned as the confidence score to that boundary. This is repeated for all boundaries in all images. Due to noise in the seismic data and the coarseness of this method, false positives were a problem. Because the goal is a regional map and density of the data is generally high, the lowest 50% of the CIT scored-boundaries were removed from mapping. This value was determined based on the examination of results using different thresholds. The CIT scores for the remaining boundaries were divided by quartile for grades: “A” for the top 25% of CIT scores, “B” for the next 25% down, “C” for the second to lowest class, and “D” for the bottom quartile. On the basis of visual inspection and mindful of the age of the seismic used in the study, boundaries estimated at depths greater than 6 s were dropped.

Depth conversion and mapping

Although the results of the texture and reflector analysis estimate entire salt bodies, only the top of the salt is of interest here. The time to bottom of salt was also examined, but in many of the seismic sections the bottom of salt was not evident even by manual inspection. The time to top of salt was extracted from the estimated salt boundaries in each line and translated vertically to sit along the navigation line. These points were then converted to depth based on a velocity field estimated from 3593 velocity surveys. These surveys were submitted to the U.S. government by operators and are publicly available (USBOEM, 2018). Although it is typical for velocity to increase with depth, a constant velocity was assumed and linear regressions of time versus depth were estimated. The average goodness-of-fit statistic (R^2) for the regressions was 0.99, which justified using that functional form and its parameters as a good first-order approximation of the local velocity field. The 3593 estimates of average velocity were gridded for a Gulf-wide velocity map using an anisotropic ordinary kriging model (ESRI, 2018). These estimated velocities were used to convert the time to top of salt to depth for every point in the study area.

Because the original data from the study included dozens of separately shot surveys, in some locations, multiple lines of different origins were in very close proximity. Because the analysis of each line was performed independently, a single location (or a very small neighborhood around a single location) could have multiple depths to the top of salt. If inverse distance weighting (IDW) gridding was directly applied to the union of all points, the result would be an extremely high variance surface that was much more a reflection of the diversity of the data rather than the depth of the top of salt.

Therefore, the following algorithm was adopted. First, a grid of base points, 500 m (1640 ft) apart was established over the study area. Around each base point, all points representing the depth to the top of salt within 350 m (1148 ft) of each point were found and grouped by their line number. Then, all observations from a single line in that group were averaged. Those line-specific average depths in that neighborhood that were more than 25% deeper than the shallowest average were dropped from the calculation. In total, 21,716 points were eliminated from the 241,416 points. Then, IDW was applied, with a power of 1, to estimate the depth to top of salt surface (ESRI, 2018). The final clipping rule was applied so that no part of the estimated surface extended more than 2000 m (6560 ft) from the nearest control.

COMPUTATIONAL COMPLEXITY

On a single desktop computer, running an Intel i5-3450, 3.10 GHz, 64 bit processor, sequential processing of 8043 seismic lines in this study required approximately 645 h. Image analysis occupied approximately 62% of this time, 31% for the reflector analysis, and 7% for CIT, input/output, and miscellaneous tasks. Our computer configuration, whereas not *low end*, was specifically kept to a minimalist configuration in the spirit of measuring computational complexity for the worst case.

The computational complexity of processes in this workflow, for each line, was linear in the seismic image dimensions (mn , where m is the length of the image, n is the width, and mn is the total number of pixels of the image). In the case of the GLCM, the process had a computational complexity of $O(N^2mn)$ (where N is the length of the search neighborhood); whereas this expression is linear in mn , it is not in N , which greatly increases computation time for larger search neighborhoods in the GLCM.

Because of the analytic independence of processing each line, a linear gain in computation time can be made by parallelizing the project across multiple processors. Likewise, because of the centrality of matrix manipulation in our approach (particularly, GLCM), switching from a standard central processing unit to a graphics processing unit (GPU)-technology-based hardware platform also seems to be a very effective step in accelerating the analysis (perhaps as much as 10×). Therefore, a 10-processor-GPU array could shrink computer time by a factor of as much as 50–100 — reducing run time from a little less than a month to, perhaps, less than three days.

RESULTS AND DISCUSSION

The principal result of the project is shown in Figure 8, a basin-wide depth to the top of the salt map. From it, we are in a position to judge the present research on the four criteria set forth at the outset.

First, did the workflow and its components yield technically useful results? The study area in Figure 8 is 278,000 km² (107,000 mi²). Of the total study area, 17% had insufficient seismic coverage for our analysis. We removed 16% of the portion that we *could* map from the final products for failure to meet our quality control thresholds. Qualitatively, we compared the depth to salt by inspection of individual lines across the basin. Misses occurred, of course. However, the final maps meet the needs of regional to basin-wide geologic mapping — one does not plan a well on a basin-wide map. The practical utility of our final mapping was dramatically improved by the systematic reporting and use of assessment of certainty in the analysis (Figure 9).

Our top of salt compared well with published regional and basin-wide work (Fort and Brun, 2012; Ganey-Curry, 2018). A larger scale assessment of accuracy on the shelf is shown in two comparisons in Figure 10. In that map, which focuses on the central part of the GOM shelf, south of Louisiana, our estimated shallow and intermediate diapiric salt bodies are shown in deep green. Compared to these are the results of a project in the mid-1990s to hand-map shallow salt across the GOM shelf (Hentz et al., 1997). The shallow/diapiric salt that they found is symbolized in Figure 10 by purple polygons. Finally, Figure 10 includes field outlines for oil and gas fields identified as trapped by intermediate to shallow salt. These classifications are made by the U.S. government on the basis of data submitted to them by the operators (USBOEM, 2018). The outlines of those fields appear in Figure 10 as blue striped polygons. Our mapping of intermediate to shallow salt in this area has a very high spatial correlation with the results of *both* independent studies.

These comparisons lend support to the validity of our mapping of the *top* of the salt and do not provide guidance on the feasibility of




Figure 8. The final depth to top of salt map for the study area within the U.S. portion of the GOM. The dark black line shows the limits of the study area. White areas within the study limits were not mapped either because there were no data (or insufficiently dense data) within them, or the confidence in the estimated depth to the top of salt did not meet the quality control thresholds. The gray boundaries in the map represent U.S. government protraction areas (which organize leasing).

mapping salt thickness or the *roots* of salt structures, particularly at substantial depths (which motivated the decision not to map below 6 s). As well, while taking effective steps to reduce false positives in the analysis (mapping salt where there appears to be none), at this stage of the research, we still do not have a scale-robust method for identifying false negatives (mapping no salt on a line where salt

appears on visual examination). Very small holes (i.e., nonsalt) in the middle of large salt bodies are removed by morphological editing. However, a method of automatically identifying a seismic line with *no* salt is an important goal that has not been achieved yet.

Second, is the approach scalable? Although we made methodological additions and amendments, most of the individual salt discrimination methods used were previously introduced in the geophysical literature, originally using a very small number of lines or synthetic data. Testing the limits of their scale of application (as modified by us) was a fundamental project goal. At the basin scale, the quality of our results would have suffered had we not developed the process of macroediting, or looking for errors on a survey or regional scale.

Nevertheless, at least through a project of 8000 + lines, we encountered no obvious limitation on the volume of input data susceptible to a workflow of batch, automated image processing. In addition, with respect to scale, the practical utility of the study to working geoscientists without an included analysis of uncertainty shown in Figure 9 would have been lower. With the basin-wide certainty mapping, as well as assessments of every individual boundary, a geologist or geophysicist can zoom in on his or her own project area and automatically see consistent evaluation of the reliability of this work.

Third, the analysis appeared robust in the face of the diversity of 82 input surveys, acquired between

1981 and 1992. The macroediting procedure specifically attacked survey-level problems, but with that tool in hand, the effect of variance in the input data appears to have been minimized (or at least, significantly reduced). Our macroediting approach, however, while correcting regional errors and improving accuracy, does not change the *regional* nature of our mapping or make this a high-precision estimate of boundaries (especially on the flanks of salt structures).

Of course, accuracy would be immeasurably improved had we based the project on the newest and most technically homogeneous data. There would have also been a smaller area removed for quality control reasons. However, there is nothing in the methodology that *requires* old or diverse data — it will all work much better with the *latest and greatest* inputs. Nevertheless, using *this* data set tested the utility of a machine learning approach to exploit the *massive idle stores of older data*, and that goal was accomplished with a positive outcome. The fact that it will work with superior data is just a (very large) upside.

The final dimension of performance was economic: Does the value derived from batch, automated image analysis, such as presented here, exceed the marginal cost of conducting it? Ultimately, this question must be left to companies holding seismic data resources similar to what we used (or better). However, in answering this question, realize that the original costs of acquisition and processing of these thousands of lines were sunk long ago. Moreover, although further research and development of our approach is definitely warranted, its forward costs are relatively low. Therefore, if this project is representative, it is hard to imagine that the economic benefit of this approach does not exceed by many times the costs of its application.




Figure 9. The final confidence grades assessed for each estimated salt/nonsalt boundary. The assessments were restricted to the top 25% of all boundaries based on their length (eliminating small boundaries that were likely processing effects) and the top half of certainty scores, based on CIT.




Figure 10. The correspondence on the shelf between salt domes estimated in this study (dark green) and two published studies. The blue striped polygons are the outlines for fields identified with shallow salt trapping (USBOEM, 2018). The purple polygons are hand-mapped shallow salt (Hentz et al. 1997). Latitude and longitude coordinates are provided for reference, and the labeled protrusion areas are South Marsh Island (SM); Eugene Island (EI); and Ship Shoal (SS).

CONCLUSION

Salt is central to the habitat of hydrocarbons in the GOM and a host of highly petroliferous basins around the world. Small-scale studies had shown that salt could be located automatically, but our research extended those findings in the scale, volume, and difficulty of input data quality. These successes, in turn, open the door to potentially very profitable reuse of massive vintage seismic resources that today mainly occupy storage space.

When we move from the shelf to the slope, there are still extensive areas where the certainty in the analysis remains high. Indeed, there are some areas on the slope, where the horizontal, canopy geometries dominate, in which the identification of salt appears very good. Because salt bodies are thinner, there are even some cases in which the top and bottom of salt are clearly distinguished from the surrounding rock.

However, on the slope the results also show that spatial variance in the assessment of certainty of results is higher than on the shelf. That is, on the shelf, there are broad and continuous areas receiving the same certainty score. On the slope, whereas some areas have high certainty scores, they tend to be smaller in area and less continuous. The variable morphology of salt on the shelf probably drives this difference and suggests that experiments with different models between these geologically disparate provinces are justified in future research.

On a more general level, it must also be admitted that, in its classic embodiment as diapirs on a wide, passive margin shelf, salt is one of the easier targets for pattern recognition. More nuanced features will pose stiffer challenges: Faults, facies, rock properties, and fluid types are all critically important. Features that synthesize empirical geophysical properties and geologic theory, such as facies, will probably raise the greatest challenges. Yet without success using these techniques, applied to salt on an industrial scale, the motivation to move on to more complex tasks would be undermined.

In testing a new technology, its limitations must be as clear as its advantages. Most broadly, nothing from our experience supports the popular notion that data may be dumped into an algorithm and useful answers will magically appear. Expertise and many experiments were required to design the workflow, find efficient and reliable discriminators, test the results, and respond to systematic errors found. Moreover, irrespective of the provenance of input data, when scale exceeds the limit of hand checking output, automated quality control assessment *must* accompany the results for them to be accepted and put to practical use.

This research demonstrated the value of batched, automated analysis to extract valuable regional geologic information from a massive and diverse collection of geophysical data. Equally important: The component tools we applied are broadly accessible to those who sit on the massive volumes of data to which they can be applied. After all, it is more experience in application that will best distill the technical and economic value of automated analysis of massive data sets from the lustrous aura that surrounds it.

ACKNOWLEDGMENTS

The authors express their appreciation to T. Cullison, J. Castagna, and three anonymous reviewers for their suggestions to improve the present paper. We also wish to thank Earth Science Associates for their support of the research and permission to use their GOM³ software for the data and analysis described here.

DATA AND CODE AVAILABILITY

All data used in this study are publicly available. The raw SEG-Y and associated SEG-P1 seismic data are available from the USBOEM via a website of the U.S. Geological Survey at <https://walrus.wr.usgs.gov/namss/search/> (Triezenberg et al., 2016). The data for velocity surveys and geopolitical/regulatory data (e.g., lease blocks, protraction areas, and state boundaries) are available through a website of the USBOEM, <https://www.data.boem.gov/>. Although all raw data used here are available from public sources, the versions of data sets used were all drawn from the Earth Science Associates GOM³ system. The code developed for this study is proprietary to Earth Science Associates.

DATA AND MATERIALS AVAILABILITY

Data associated with this research are available and can be accessed via the following URL: doi: [10.5066/F7930R7P](https://doi.org/10.5066/F7930R7P).

APPENDIX A

GLCM

Using GLCM consists of two parts:

- 1) Build the GLCMs for all pixels: Establish an examination pixel near the upper left corner of the image, and define a neighborhood of $N \times N$ pixels centered on it (here, $N = 51$). (This discussion

Table A-1. GLCM statistics used.³

GLCM statistic	Equation	Measures
Contrast	$\alpha_{mn}^1 = \sum_{ij=1}^G p_{ij}(i-j)^2$	Amount of gray-level variation in search neighborhood.
Homogeneity	$\alpha_{mn}^2 = \sum_{ij=1}^G \left(\frac{p_{ij}}{1+(i-j)^2} \right)$	The smoothness of the gray-level distribution in the search neighborhood.
Dissimilarity	$\alpha_{mn}^3 = \sum_{ij=1}^G p_{ij}(i-j)$	Similar to contrast but with a smaller weight on p_{ij} .
Entropy	$\alpha_{mn}^4 = \sum_{i=1}^G \sum_{j=1}^G p_{ij} \ln p_{ij}$	Degree of disorder in the gray levels of the pixels in the search neighborhood.

Note: α_{mn}^k = the k th GLCM statistic, estimated for a target pixel located at m, n in the seismic image. m, n = row and column indices of cells (pixels) in the entire seismic line image; i, j = row and column indices for cells in the GLCM; G = the number of levels in the gray scale; p_{ij} = empirical marginal probability of co-occurrence of the i th and j th gray levels within the search neighborhood (i.e., the empirical counts of gray-level i horizontally adjacent to gray-level j for the search neighborhood, divided by the total number of possible co-occurrences ($N \times (N - 1)$)).




Figure A-1. The process for calculation of the GLCM for a single pixel within a seismic line. (a) A piece of the study seismic line showing identified target pixel for the GLCM analysis (black dot) and a search neighborhood surrounding it. (b) Magnified view of the area within the black box in (a) showing identified piece of a row (black rectangle) along which the co-occurrences are counted row-wise between horizontally adjacent cells (pixels). (c) Two horizontally adjacent pixels that are outlined — the left cell has a grayscale value of 4, and the adjacent cell to the right has a grayscale value of 1. (d) The single co-occurrence counted in (c) recorded as an entry of 1 in the fourth row and first column. (e) The GLCM after all pairs of horizontally adjacent cells are counted and all counts in the cells of the matrix are divided by the number of observations in the search neighborhood (in our case, 2550 (50 × 51)) to produce a matrix of marginal probabilities of the co-occurrence of all possible combinations of grayscale values.

ignores the treatment of edge effects, which satisfies the sampling requirements of creating marginal probabilities in the last step of creating the GLCM.) Then, starting in the upper left corner of the neighborhood, count the number of co-occurrences of gray levels in successive pairs of pixels on the top row of the neighborhood, moving right; then, drop to the row below and repeat until all rows (and pixel pairs) in the neighborhood are accounted. The process is illustrated in Figure A-1. (It is not the only GLCM estimable for a pixel because the algorithm can be run horizontally, vertically, or diagonally — horizontally either left-right or vice versa; alternatively, in the vertical and diagonal cases, up-down or vice versa).

For a system of G gray levels, this creates a $G \times G$ matrix of the empirical counts of gray-level co-occurrences for the neighborhood of one pixel (here, $G = 7$). Normalizing the count matrix by dividing each element by the total number of possible outcomes ($N \times (N - 1) = 2550$) produces a 7×7 matrix of the joint marginal probabilities of gray-level co-occurrences (p_{ij}) between the seven gray levels (i and j) among horizontally adjacent pairs of pixels, examined row-wise left to right. This is the GLCM for a pixel.

- 2) Compute statistics for each pixel's GLCM: Several statistics have been applied to the GLCM to characterize the texture of a neighborhood, as described by Hall-Beyer (2017), Albregtsen (2008), and Haralick et al. (1973). Generally, they fall into three classes: (1) probability distribution descriptors: mean, variance and energy, (2) linear correlation between the gray-level values of pairs of adjacent pixels, and (3) sums of the GLCM cell values (p_{ij}),

with each cell weighted by functions of the distance between the cell location in the GLCM and its principal diagonal (which contain the G marginal probabilities of co-occurrences of the same gray-level. The GLCM library in R automatically computes eight statistics.

REFERENCES

- Albregtsen, F., 2008, Statistical texture measures computed from gray level cooccurrence matrices: Department of Informatics, University of Oslo, 1–14, <https://www.uio.no/studier/emner/matnat/ifi/INF4300/h08/undervisningsmateriale/gldcm.pdf>, accessed 20 July 2018.
- Amin, A., M. Direche, T. Hegazy, Z. Wang, and G. Al-Regib, 2015, A novel approach for salt dome detection using a dictionary-based classifier: 85th Annual International Meeting, SEG, Expanded Abstracts, 1816–1820, doi: [10.1190/segam2015-5925748.1](https://doi.org/10.1190/segam2015-5925748.1).
- Araya-Polo, M., T. Dahlke, C. Frogner, C. Zhang, T. Poggio, and D. Hohl, 2017, Automated fault detection without seismic processing: The Leading Edge, **36**, 208–214, doi: [10.1190/le36030208.1](https://doi.org/10.1190/le36030208.1).
- Baraldi, A., and F. Parmiggiani, 1995, An investigation of textural characteristics associated with gray level cooccurrence matrix statistical parameters: IEEE Transactions on Geoscience and Remote Sensing, **33**, 293–304, doi: [10.1016/S0034-4257\(96\)00156-3](https://doi.org/10.1016/S0034-4257(96)00156-3).
- Berthelot, A., A. H. Solberg, and J.-L. Gelius, 2013, Texture attributes for detection of salt: Journal of Applied Geophysics, **88**, 52–69, doi: [10.1016/j.jappgeo.2012.09.006](https://doi.org/10.1016/j.jappgeo.2012.09.006).
- Bhosle, V. V., and V. P. Pawar, 2013, Texture segmentation: Different methods: International Journal of Soft Computing and Engineering, **3**, 69–74.
- CDA, 2018, www.ukoilandgasdata.com, accessed 15 May 2018.
- ESRI, 2018, ArcGIS/ArcGIS-Pro documentation, Redlands, CA. For the Raster to Polygon Tool: <http://pro.arcgis.com/en/pro-app/tool-reference/conversion/raster-to-polygon.htm>; for kriging and inverse distance weighting grid estimation: <https://www.esri.com/en-us/arcgis/products/geostatistical-analyst/overview>; for Anselin Local Moran's I tool, http://resources.esri.com/help/9.3/arcgisengine/java/gp_toolref/spatial_statistics_tools/cluster_and_outlier_analysis_colon_anselin_local_moran_s_i_spatial_statistics_.htm, accessed 21 June 2018.
- Fort, X., and J.-P. Brun, 2012, Kinematics of regional salt flow in the northern Gulf of Mexico: Geological Society, London, Special Publications **363**, 265–287.
- Ganey-Curry, P. E., 2018, Gulf basin depositional synthesis (GBDS) phase XI year 2 atlas: University of Texas, Institute of Geophysics.
- Gao, D., 2003, Volume texture extraction for 3D seismic visualization and interpretation: Geophysics, **68**, 1294–1302, doi: [10.1190/1.1598122](https://doi.org/10.1190/1.1598122).
- Gao, D., 2011, Latest developments in seismic texture analysis for subsurface structure, facies, and reservoir characterization: A review: Geophysics, **76**, no. 2, W1–W13, doi: [10.1190/1.3553479](https://doi.org/10.1190/1.3553479).
- Hall-Beyer, M., 2017, GLCM texture: A tutorial: Department of Geography, University of Calgary, 3, https://prism.ucalgary.ca/bitstream/handle/1880/51900/texture%20tutorial%20v%203_0%20180206.pdf?sequence=11&isAllowed=y, accessed 15 May 2018.
- Halpert, A., and R. G. Clapp, 2008, Salt body segmentation with dip and frequency attributes: Stanford Exploration Project, 113, <http://sepwww.stanford.edu/data/media/public/docs/sep136/adam1/paper.pdf>, accessed 18 May 2018.
- Haralick, R. M., K. Shanmugam, and I. Dinstein, 1973, Textural features for image classification: IEEE Transactions on Systems, Man, and Cybernetics, **SMC-3**, 610–621, doi: [10.1109/TSMC.1973.4309314](https://doi.org/10.1109/TSMC.1973.4309314).
- Hegazy, T., and G. Alregib, 2014, Texture attributes for detection salt bodies in seismic data: 84th Annual International Meeting, SEG, Expanded Abstracts, 1455–1459, doi: [10.1190/segam2014-1512.1](https://doi.org/10.1190/segam2014-1512.1).
- Hegazy, T., Z. Wang, and A. Ghassan, 2015, The role of perceptual texture dissimilarity in automating seismic data interpretation: IEEE Global Conference on Signal and Information Processing, doi: [10.1109/GlobalSIP.2015.7418168](https://doi.org/10.1109/GlobalSIP.2015.7418168).
- Hentz, T. F., S. J. Seni, and E. G. Wermund Jr., 1997, Atlas of northern Gulf of Mexico gas and oil reservoirs: The University of Texas at Austin, Bureau of Economic Geology.

- Otsu, N., 1979, A threshold selection method from gray-level histograms: IEEE Transactions on Systems, Man, and Cybernetics, **SMC-9**, 62–66, doi: [10.1109/TSMC.1979.4310076](https://doi.org/10.1109/TSMC.1979.4310076).
- Pau, G., F. Fuchs, O. Sklyar, M. Boutros, and W. Huber, 2010, EBImage — An R package for image processing with applications to cellular phenotypes: Bioinformatics, **26**, 979–981, doi: [10.1093/bioinformatics/btq046](https://doi.org/10.1093/bioinformatics/btq046).
- R Core Team, 2013, R: A language and environment for statistical computing, [www.R-Project.org](http://www.R-project.org), accessed 15 May 2018.
- Shafiq, M. A., Y. Alaudah, H. Di, and G. AlRegib, 2017a, Salt dome detection within migrated seismic volumes using phase congruency: 87th Annual International Meeting, SEG, Expanded Abstracts, 2360–2365, doi: [10.1190/segam2017-17443246.1](https://doi.org/10.1190/segam2017-17443246.1).
- Shafiq, M. A., Z. Wang, G. AlRegib, A. Amin, and M. Deriche, 2017b, A texture-based interpretation workflow with application to delineating salt domes: Interpretation, **5**, no. 3, SJ1–SJ19, doi: [10.1190/INT-2016-0043.1](https://doi.org/10.1190/INT-2016-0043.1).
- Shafiq, M. A., Z. Wang, A. Amin, T. Hegazy, M. Deriche, and G. AlRegib, 2015, Detection of salt-dome boundary surfaces in migrated seismic volumes using gradient textures: 85th Annual International Meeting, SEG, Expanded Abstracts, 1811–1815, doi: [10.1190/segam2015-5927230.1](https://doi.org/10.1190/segam2015-5927230.1).
- Triezenberg, P. J., P. Hart, and J. Childs, 2016, National Archive of Marine Seismic Surveys (NAMSS): A USGS data website of marine seismic reflection data within the U.S. Exclusive Economic Zone (EEZ): U.S. Geological Survey Data Release, <https://www.doi.org/10.5066/F7930R7P>, accessed 21 June 2018.
- United States Bureau of Ocean Energy Management, 2018, Original data available through their public website, <https://www.data.boem.gov/Main/Default.aspx>, accessed 21 June 2018.
- Urbach, E. R., and M. H. F. Wilkinson, 2008, Efficient 2-D grayscale morphological transformations with arbitrary flat structuring elements: IEEE Transactions on Image Processing, **17**, 1–8, doi: [10.1109/TIP.2007.912582](https://doi.org/10.1109/TIP.2007.912582).
- Wallet, B. C., and R. E. F. Pepper, 2013, Using mathematical morphology in an attribute workflow to improve the interpretability of salt bodies in the Gulf of Mexico: 83rd Annual International Meeting, SEG, Expanded Abstracts, 1324–1328, doi: [10.1190/segam2013-0941.1](https://doi.org/10.1190/segam2013-0941.1).
- Wang, Z. T., T. Hegazy, Z. Long, and G. AlRegib, 2015, Noise-robust detection and tracking of salt domes in postmigrated volumes using texture, tensors and subspace learning: Geophysics, **80**, no. 6, WD101–WD111, doi: [10.1190/geo2015-0116.1](https://doi.org/10.1190/geo2015-0116.1).
- Wu, X., 2016, Methods to compute salt likelihoods and extract salt boundaries from 3D seismic images: Geophysics, **81**, no. 6, IM119–IM126, doi: [10.1190/geo2016-0250.1](https://doi.org/10.1190/geo2016-0250.1).
- Zvovoff, A., 2016, Calculate textures from grey-level co-occurrence matrices (GLCMs): Package “glcm,” <https://cran.r-project.org/web/packages/glcm/glcm.pdf>, accessed 20 June 2018.