Geofísica Internacional (2025) **64-4**: 1803- 1814

GLCM and k-means clustering for texture-based recognition of salt bodies in seismic images

M.G. Orozco del Castillo[1], J.E. Abreu-Torres[2], J.J. Hernández-Gómez[3] and J.C. Ortiz-Alemán[4]

Abstract

Subsalt imaging and interpretation remain challenging tasks for oil and gas exploration, as well as for identifying suitable sites for storing large amounts of atmospheric carbon dioxide. The complex geometry of salt bodies and the poor visibility of subsalt structures introduce significant uncertainty into manual seismic interpretation, especially when defining the base of salt formations. To address this, we present an interpretable and lightweight semiautomatic methodology for salt body delineation in 2D seismic sections. This approach combines texture features derived from the Gray Level Co-occurrence Matrix (GLCM) with k-means clustering. We performed a comprehensive parameter sweep, exploring different window sizes and cluster counts, evaluated using strict consensus labels obtained from two expert interpretations. A 5-fold cross-validation procedure was applied to identify configurations with consistently high *F1* scores, achieving a maximum *F1* of 0.78. To enhance robustness, we introduced a consensus-based strategy that aggregates results from multiple high-performing configurations. The resulting consensus map provides a continuous confidence measure for salt likelihood, highlighting well-defined boundaries and subtle features often overlooked in manual interpretations. This method reinforces the hypothesis that salt bodies exhibit a distinctive textural signature in the GLCM feature space and demonstrates the potential of interpretable unsupervised learning to support expert workflows in seismic interpretation involving salt tectonics.

Key words: Seismic interpretation, Subsalt exploration, Texture-based recognition, Unsupervised learning, Gray Level Co-occurrence Matrix.

Resumen

La obtención de imágenes e interpretación de estructuras subsalinas sigue siendo una tarea compleja para la exploración de petróleo y gas, así como para la identificación de sitios adecuados para almacenar grandes cantidades de dióxido de carbono atmosférico. La geometría irregular de los cuerpos salinos y la baja visibilidad de las estructuras bajo la sal generan gran incertidumbre en la interpretación sísmica manual, en especial al definir la base de las formaciones salinas. Para abordar este desafío, presentamos una metodología semiautomática, interpretable y ligera para la delimitación de cuerpos salinos en secciones sísmicas 2D. Esta combina características de textura derivadas de la matriz de coocurrencia de niveles de gris (GLCM) con agrupamiento mediante k-means. Realizamos un análisis exhaustivo de parámetros, explorando distintos tamaños de ventana y números de clústeres, evaluados con etiquetas de consenso derivadas de interpretaciones de dos expertos. Mediante validación cruzada de 5 pliegues, identificamos configuraciones con puntuaciones *F1* consistentemente altas, alcanzando un máximo de 0.78. Para mayor robustez, integramos una estrategia de consenso que combina resultados de configuraciones destacadas. El mapa de consenso resultante ofrece una medida continua de confianza sobre la probabilidad de presencia de sal, destacando límites claros y rasgos sutiles que suelen pasarse por alto en las interpretaciones manuales. Este enfoque refuerza la hipótesis de que los cuerpos salinos poseen una firma textural distintiva en el espacio de características GLCM y muestra el potencial del aprendizaje no supervisado interpretable para complementar el trabajo de los expertos en tectónica salina.

Palabras clave: Interpretación sísmica, exploración subsalina, reconocimiento basado en texturas, aprendizaje no supervisado, matriz de coocurrencia de niveles de gris.

Received: October 23, 2024; Accepted: July 10, 2025; Published on-line: October 1, 2025.

Editorial responsibility: Dr. Joel Rosales-Rodríguez

* Corresponding author: Mauricio Gabriel Orozco del Castillo, mauricio.orozco@itmerida.edu.mx
[1] Tecnológico Nacional de México/IT de Mérida, Departamento de Sistemas y Computación, Mérida, Yucatán, México
[2] King Abdullah University of Science and Technology (KAUST), Ali I. Al-Naimi Petroleum Engineering Research Center, Physical Sciences and Engineering Division, Thuwal, Saudi Arabia
[3] Instituto Politécnico Nacional, Centro de Desarrollo Aeroespacial, Ciudad de México, México
[4] Centro de Investigación Científica de Yucatán, Unidad de Ciencias del Agua, Unidad Cancún 77524, Cancún, Quintana Roo, México

Orozco del Castillo, M. G.; Abreu-Torres, J. E.; Hernández-Gómez, J. J.; Ortiz-Alemán, J. C.

https://doi.org/10.22201/igeof.2954436xe.2025.64.4.1833

1. Introduction

Seismic imaging within salt tectonic regions presents a formidable challenge due to the diminishing resolution as we go deeper beneath the overlying salt structures (Jones & Davison, 2014). Given that potential oil reserves are often situated in close proximity to or beneath salt formations (Ratcliff *et al.*, 1992), there has been a growing interest in recent years to develop new computational tools which aim to assist seismic interpreters in estimating characteristics such as geometry, location, and depth of salt diapirs (Babakhin *et al.*, 2019; Berthelot *et al.*, 2011; Farrokhnia *et al.*, 2018; Ji *et al.*, 2020; Orozco-del-Castillo *et al.*, 2011; Orozco-del-Castillo *et al.*, 2017). Traditional processing and interpretation of seismic data can provide an appropriate location of the top of salt bodies, but accurate estimation of the irregular geometry of salt bodies, and especially the location and shape of its base, remains a difficult task for experts. Analysis of seismic volume data has generally been done through successive interpretation of adjacent seismic profiles (2D images) where interpretation criteria are similar, a task that is considered both time-consuming and tedious (Cohen *et al.*, 2006). As a consequence, in many cases only some portions of the entire seismic volume are interpreted. As the volume and resolution of seismic data increase as a consequence of evolution on seismic acquisition technology, the ability to perform semi-automated interpretations will become increasingly attractive to ensure the productivity of interpreters and minimize the subjectivity of their work (Babakhin *et al.*, 2019; Farrokhnia *et al.*, 2018; Ji *et al.*, 2020).

Seismic interpretation involves identification of characteristic features expressed as image patterns in order to subsequently apply image segmentation. This process is highly dependent on the human ability to visualize complex images. In this sense, different automatic and semiautomatic methods have been proposed to recognize particular characteristics in seismic patterns, allowing the subsequent segmentation or delimitation of images.

Over the decades, various methods have been proposed for the detection of salt bodies in seismic data, ranging from early refraction-based techniques (Musgrave, 1960) and bandpass filtering (Ng *et al.*, 1993) to more advanced textural and pattern recognition methods. Textural analysis has played a central role (Gao, 2004), including expert systems (Love & Simaan, 1985), neural networks for seismic facies mapping (West *et al.*, 2002), and Gray Level Co-occurrence Matrix (GLCM)-based descriptors (Chopra & Alexeev, 2006). With increased computational capacity, more recent approaches have integrated frequency-based attributes (Berthelot *et al.*, 2013), genetic algorithms (Couder-Castañeda *et al.*, 2025; Orozco-del-Castillo *et al.*, 2017), and deep learning techniques such as U-Net (Ji *et al.*, 2020) and ensemble convolutional neural networks (CNNs) (Babakhin *et al.*, 2019). Unsupervised learning has also emerged as a practical alternative, exemplified by Di *et al.* (2018), who used multi-attribute k-means clustering to delineate salt boundaries in 3D seismic data.

In this work, we present a semiautomatic method for recognizing salt bodies in seismic images by combining the k-means clustering algorithm with GLCMs to extract four key texture attributes: contrast, correlation, energy, and homogeneity. To ensure robustness, the method incorporates a k-fold cross-validation scheme and evaluates clustering performance across multiple random initializations. Key configurations are selected based on composite performance metrics (including *F1* score, Dice coefficient, and IoU), and their outputs are aggregated into a consensus heatmap that reflects the probability of each pixel being part of a salt body. The results are compared with expert interpretations to validate geologic plausibility. While the method is not intended to replace expert judgment, it offers an interpretable, lightweight, and computationally efficient support tool to aid geoscientists in the early stages of seismic analysis—particularly by highlighting textural patterns, exposing ambiguous regions, and potentially reducing manual workload.

This work is organized as follows. Section 2 reviews relevant literature on automatic and semiautomatic methods for salt body identification in seismic data. In Section 3, we describe the theoretical foundations of the proposed methodology, including the use of GLCM-based texture attributes and k-means clustering, as well as the characteristics of the seismic section and expert interpretations used for evaluation. Section 4 details the development of the consensus-based approach, the experimental design—including cross-validation and performance metrics—and the main results. Section 5 presents a discussion of the key findings, methodological limitations, and potential directions for future research. Finally, Section 6 provides the main conclusions and reflects on the practical implications of the study.

2. Related work

Over the years, various techniques have been developed to aid in the identification and segmentation of salt structures within seismic data. Early efforts focused on manual and semi-automated approaches, including expert systems (Love & Simaan, 1985), bandpass filtering (Ng *et al.*, 1993), and amplitude-based segmentation (Lomask *et al.*, 2004). As computational power and data availability increased, texture-based methods became more prominent, such as those employing GLCMs (Chopra & Alexeev, 2006), frequency and similarity attributes (Berthelot *et al.*, 2013), and evolutionary algorithms (Couder-Castañeda *et al.*, 2025; Orozco-del-Castillo *et al.*, 2017). More recent studies have explored the integration of unsupervised learning

with geophysical descriptors to improve salt dome delineation. For instance, Tavakolizadeh and Bagheri (2022) used attribute selection combined with SOM and SVM classifiers to improve segmentation quality, showing that a reduced feature space could enhance both interpretability and performance. Similarly, Steklain *et al.* (2022) applied structure-oriented filtering before unsupervised classification to improve salt boundary preservation, emphasizing the importance of input preprocessing and the use of coherent texture features.

As machine learning matured, particularly in the realm of deep learning, new methods emerged to automate salt body segmentation. Ji et al (2020) applied CNNs to classify salt voxels, while Babakhin *et al.* (2019) employed ensemble networks to improve segmentation accuracy. These approaches have pushed the state of the art but are often reliant on large labeled datasets, significant computational resources, and intricate network architectures. To address these limitations, recent research has explored semi-supervised and unsupervised models. For example, Geng *et al.* (2022) introduced a mean-teacher framework that enhances segmentation in low-label settings, and Xu *et al.* (2025) proposed a contrastive learning-based model that generalizes across field and synthetic data. While promising, such methods still involve complex training procedures and lack the interpretability required for transparent geophysical analysis.

A more streamlined and interpretable unsupervised alternative was introduced by Di *et al.* (2018), who applied a k-means clustering algorithm to a set of seismic and texture attributes—such as semblance and energy ratio—to delineate salt body boundaries in 3D seismic volumes. While their method effectively emphasized salt contrasts and demonstrated the potential of interpretable clustering, it focused primarily on boundary delineation and required manual region initialization, which may limit its adaptability in exploratory contexts. Moreover, although k-means is an inherently transparent technique, its interpretability can diminish in high-dimensional, heterogeneous feature spaces, where it becomes difficult to assess the influence of individual attributes on the final clusters. In contrast, the method proposed in this study aims to capture both the boundaries and internal geometry of salt bodies using a compact, low-dimensional feature space composed of four GLCM-derived texture attributes—contrast, correlation, energy, and homogeneity—applied to 2D seismic sections. This design preserves interpretability and facilitates geological insight, thus making it well suited for unsupervised, early-phase interpretation scenarios.

The choice to rely on four GLCM attributes reflects a deliberate balance between descriptive capacity and interpretability. These features are well-established in capturing core aspects of texture and can be directly linked to observable geological patterns in seismic data. By maintaining a compact and transparent feature set, the method not only reduces computational overhead but also facilitates a clearer understanding of how salt bodies manifest within seismic sections. This interpretability supports expert insight, enabling interpreters to trace how specific textural characteristics relate to geological structures. Furthermore, the simplicity of the model favors integration into resource-limited environments and enhances its potential for educational or exploratory use in contexts where more opaque or complex methods may obscure geological meaning.

While deep learning approaches have recently gained popularity for salt body detection (Babakhin *et al.*, 2019; Ji *et al.*, 2020), they typically require large labeled datasets, extensive computational resources, and significant tuning. In contrast, the method presented here provides a lightweight and interpretable alternative that can operate effectively without large sets of labeled data or high-end infrastructure. Rather than replacing more complex methods, it offers a complementary strategy particularly suited for preliminary interpretation phases, low-data scenarios, or as a transparent benchmark alongside more opaque models. In this way, it contributes to a growing toolbox of techniques aimed at supporting seismic interpretation tasks under varied practical constraints.

3. Materials and methods

3.1 Gray Level Co-Occurrence Matrix (GLCM)

GLCM analysis is a process that characterizes sets of pixels of an image according to their texture. This approach has proved to be very useful in seismic image analysis since seismic facies have a different texture than the rest of the image (Cavalin & Oliveira, 2017; Hall-Beyer, 2000). The GLCM textural analysis is carried out considering the relationship between two pixels at a time: a reference and a neighbor pixel. Depending on the geometrical position of the neighbor pixel with respect to the reference one, the offset (n,m) is defined, where $n,m \in \mathbb{N} \cup \{0\}$. The offset means that the neighbor pixel is located n pixels in the x direction and m pixels in the y direction, with respect to the reference pixel. Each pixel within the image will be the reference pixel in turn, starting in the upper left corner and moving towards the lower right corner of the matrix representing the grayscale image. For instance, the pixels that are along the edge do not have a neighbor at position (1,0) so they would not be considered in this count.

The GLCM analysis uses a window of variable size to detect how many times within such window a pixel with a value A is found subsequent in the (n,m) offset to another pixel with a value B (where A and B are usually integers between 0 and 255, not necessarily different). The count for each reference pixel is accommodated within the so called GLCM. Thus, the GLCM elements indicate the number of times that a combination of gray

values was found within the window in that instance. GLCM should be symmetric for texture attribute calculations to be successful. Then, the matrix is normalized with the number of possible occurrences.

With the normalized GLCM, several textural properties yielding real values can be obtained. Among them, there is a set of independent textural properties that can be used in order to characterize texture in an image. Following Hall-Beyer (2000), we shall consider GLCMs to be square matrices of odd order. Given a GLCM p(i,j) with row and column indexes i and j, respectively, some statistical properties that can be calculated from it are:

- Contrast. A contrast intensity value between a pixel and its neighbor in the entire image. For a constant image, the contrast value is 0. This property is also known as variance or inertia:

$$|i-j|^2 p(i,j). \tag{1}$$

- Correlation. A value of how a pixel is correlated with its neighbor throughout the image. For a constant image, the result of the correlation is undefined:

$$\frac{(i-\mu_i)(j-\mu_j)p(i,j)}{\sigma_i \sigma_j}, \tag{2}$$

where μ and σ represent the mean and the standard deviation, respectively.

- Energy (or uniformity). The sum of the squared elements in the GLCM. The energy value is 1 in a constant image:

$$p(i,j)^2. \tag{3}$$

- Homogeneity. A value that measures how close the elements are to the main diagonal of the GLCM. Homogeneity results from 1 in a diagonal matrix:

$$\sum_{i,j} \frac{p(i,j)}{1+|i-j|^2}. \tag{4}$$

Further information about GLCM textural analysis can be found in the work of Hall-Beyer (2000).

3.2 Cluster analysis

The cluster analysis is a procedure that divides a set of data into groups (clusters) whose elements share certain characteristics which are useful for solving a particular problem (Tan *et al.*, 2016). This process groups data based solely on information that describes objects or the relationship between them. Although there are different types of algorithms for clustering purposes, in this work we implement the k-means algorithm, a centroid-based or distance-based algorithm, where distances are computed to assign a point to a given cluster (Wang *et al.*, 2012).

The *k*-means algorithm requires the following steps: first, the number k of clusters into which a set of M data will be divided must be determined. Next, each cluster i=1,…,k is assigned a centroid C_i, to later determine to which cluster each data that is not a centroid belongs, through the computation of its Euclidean distance to each centroid. When the data mapping is done, the algorithm recalculates the centroids for each cluster, by determining the average of the entire data set that represents each cluster. The process is repeated iteratively until the centroids no longer change (Zhang *et al.*, 2012).

A fundamental issue in machine learning is the quantitative evaluation of models to assess how well they classify data. Several metrics have been proposed for this purpose. Three of the most commonly used are precision, recall, and the *F1* score (Patterson & Gibson, 2017). Precision (also known as positive predictive value) measures how many of the predicted positives are actually correct. Mathematically, it is the ratio of true positives to the sum of true and false positives. Recall, also called sensitivity or true positive rate, measures how well the model captures actual positives, and is calculated as the ratio of true positives to the sum of true positives and false negatives. The *F1* score combines both precision and recall into a single metric using the harmonic mean:

$$F_1 = 2 * \frac{precision \cdot recall}{precision + recall} \tag{5}$$

The closer these metrics are to 1.0, the better the classification performance.

In addition to these, this study also employs two other widely used segmentation metrics: the Dice coefficient and the Intersection over Union (IoU). Both are overlap-based measures that assess how well the predicted regions align with ground truth labels. The Dice coefficient (also known as the Sørensen–Dice index) emphasizes the similarity between two sets by comparing their overlap to their combined size. IoU (also called the Jaccard index) measures the ratio of the intersection to the union of predicted and true regions. While similar in spirit to the *F1* score, both metrics are commonly used in image segmentation tasks due to their direct interpretability in terms of spatial agreement between shapes.

3.3 Seismic data

For the development of this work, images of different seismic sections with salt tectonics, previously interpreted by experts, were taken into consideration. One of such images can be observed in Figure 1. Notice the huge salt structure on the right end of the section, just below the seabed.

The seismic section on Figure 1 was independently interpreted by two different experts, and their interpretations are shown in Figure 2. With the seismic section as a grayscale image, the salt body is identified in red, and the subsalt section in blue. Notice how both interpreters define basically the same top of the salt but differ considerably with respect to the base. While Expert 1 defines a thin salt body, Expert 2 defines a much broader one. These discrepancies are not uncommon in the field of seismic exploration and interpretation, and this is one of the reasons why this is such a challenging subject. While the top of the salt body is relatively easy to interpret due to the strong reflector which identifies it, the salt base is a much more complex task. Several diffraction patterns commonly known as "smiles" are present below the salt feature, leading to disruption of seismic horizons associated with the base of salt and other subsalt events.

There is no obvious visual evidence (to the human eye, at least) of where the salt body ends. As a result, interpreters must rely heavily on their experience to produce the most accurate interpretation possible.

4. Development and results

Following the recommendations of Hall-Beyer (2000), we used unitary offsets in all eight principal directions: (0,1), (0,–1), (1,0), (–1,0), (1,1), (1,–1), (–1,1), and (–1,–1). For each valid element in the bidimensional seismic section, we computed the four classical GLCM texture properties—contrast, correlation, energy, and homogeneity—as defined in Equations (1)–(4). This resulted in a four-dimensional feature vector per pixel, yielding over one million valid elements in total (1126 traces per inline and 926 samples per trace).

We then applied k-means clustering as an unsupervised technique to group pixels based on textural similarity. Since the optimal number of clusters and texture window size were unknown, we performed a comprehensive evaluation across a grid of parameter combinations. Specifically, we tested window sizes from 3 to 61 (in steps of 2) and cluster counts from 1 to 10, resulting in 300 unique configurations. To assess the discriminative power of each configuration, we applied 5-fold cross-validation on a stratified subset of pixels sampled from the expert-labeled regions. Although

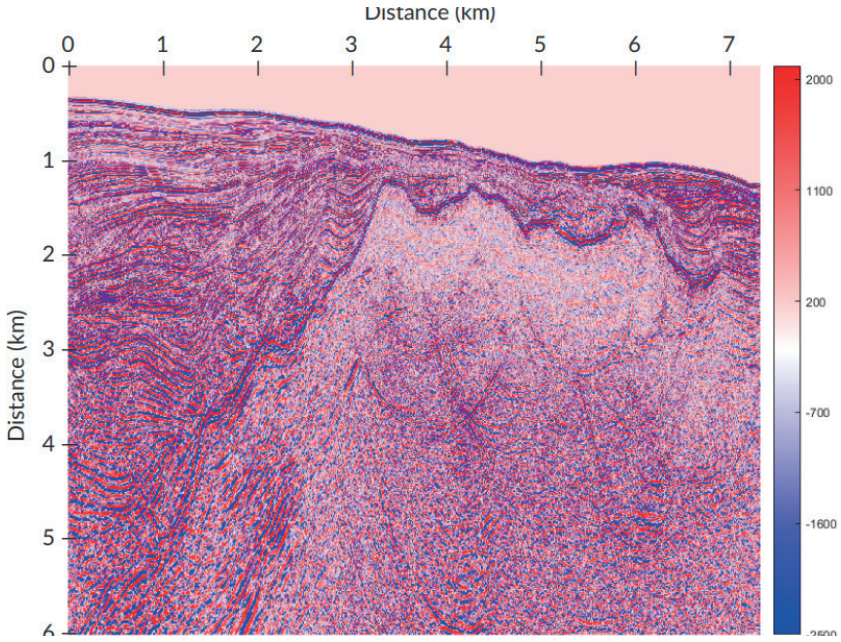

Figure 1. Pre-stack Kirchhoff depth-migrated seismic section used for interpretation and analysis. The image corresponds to the result of the third migration iteration, displayed using unfiltered amplitude values without gain adjustment. The color scale follows a red–white–blue palette, where red denotes positive amplitudes and blue denotes negative amplitudes. Vertical and horizontal axes represent distance in kilometers. This section highlights key features of the salt body, including a well-defined top reflector and steep diapir flanks. At the base of the salt dome, diffraction patterns are visible—likely caused by the limitations of the selected migration velocity (saline flooding), which impairs accurate imaging of the salt base.

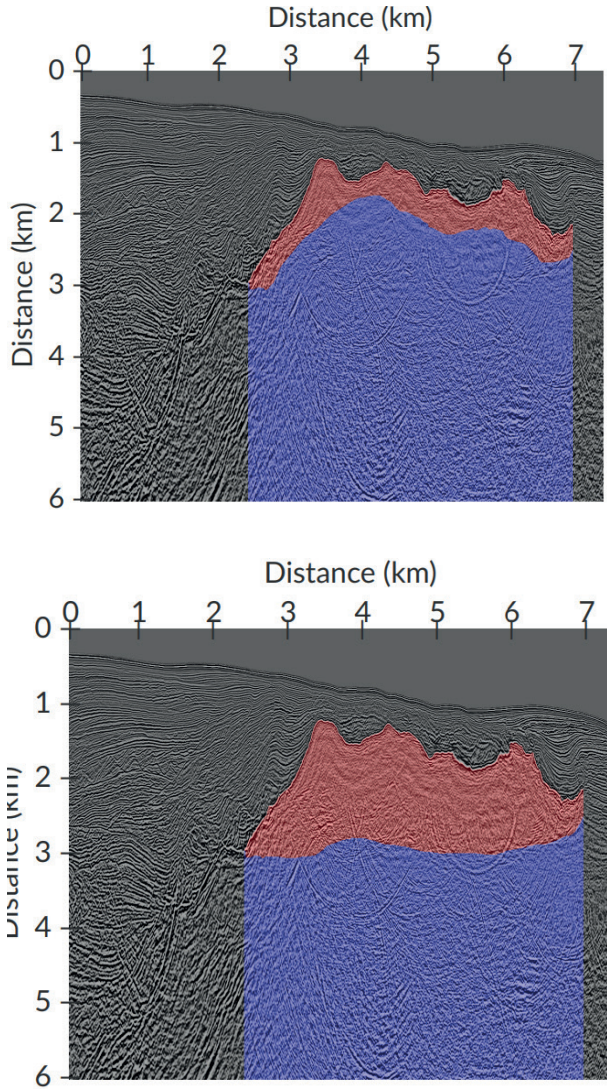

Figure 2. The interpretation of the salt body according to the human experts. The salt region is shown in red, in blue the subsalt region, while the rest of the seismic section is presented as a grayscale image. Both interpretations coincide with respect to the top of the salt body, while considerably differ with respect to the base of it.

the clustering itself is unsupervised, this evaluation framework allowed us to quantify how well each configuration could separate salt from non-salt textures. For this purpose, pixels labeled as salt by both interpreters were treated as high-confidence positives, while pixels labeled as non-salt by both were treated as high-confidence negatives. This strict labeling ensured the reliability of the evaluation subsets used for configuration scoring.

For each fold and configuration, classical texture features were extracted using GLCMs, and the training data were clustered using the k-means algorithm. To account for the randomness introduced by k-means initialization, the clustering procedure was repeated ten times per configuration, and the run yielding the highest performance on the training fold was retained for further evaluation. The cluster most predictive of salt was then used to label the test set, and performance was evaluated using multiple metrics, including accuracy, *F1* score, Dice coefficient, and IoU. To summarize performance across folds, the average value for each metric was computed. Figure 3 presents heatmaps for the *F1* score (a), Dice coefficient (b), and IoU (c), all of which consistently highlight the same high-performing region of the parameter space. These results guided the selection of the optimal window size and number of clusters for subsequent analysis. Given the consistency between these three metrics, subsequent analyses and model selection were based on the *F1* score, which balances precision and recall.

Despite the absence of a simple linear relationship between window size and number of clusters, the heatmaps in Figure 3a–c (*F1* score, Dice coefficient, and IoU) reveal a well-defined region where segmentation performance is consistently strong. Several configurations—particularly those with small to moderate

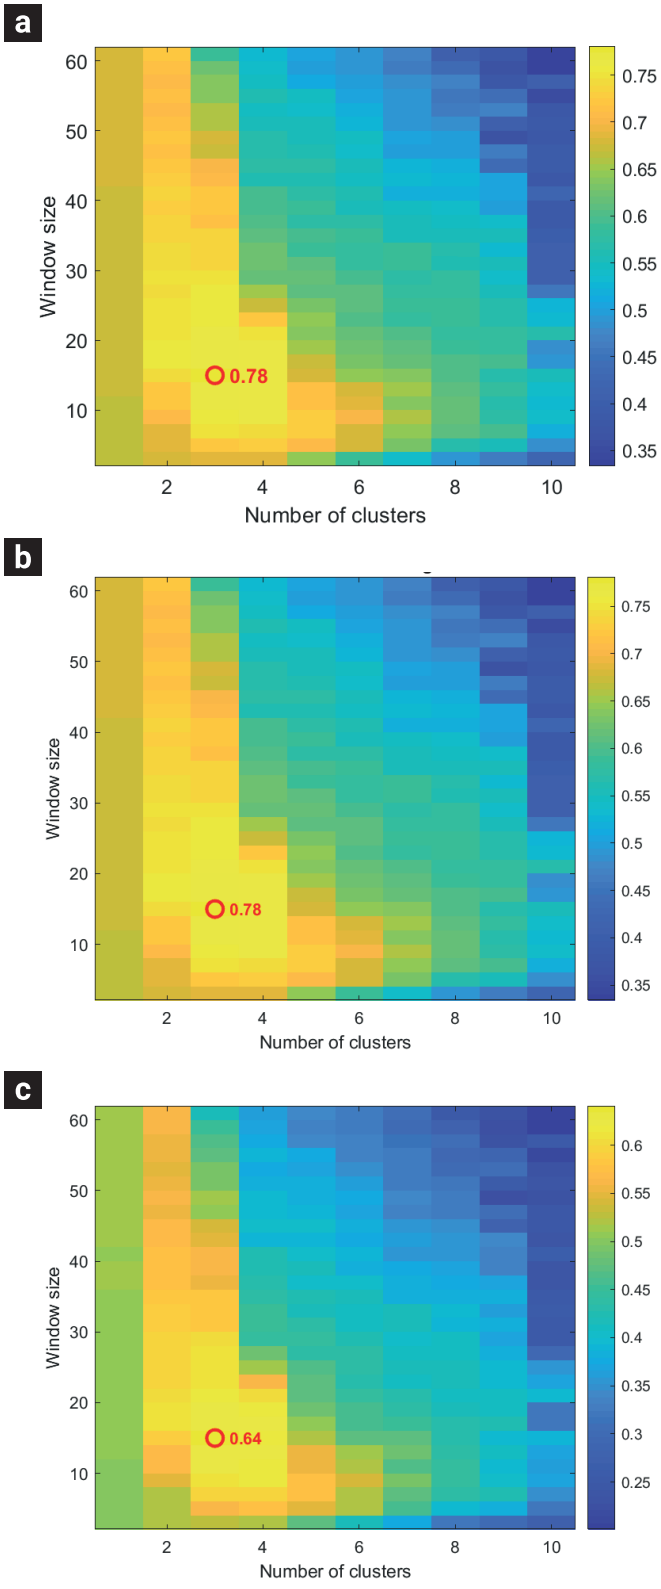

Figure 3. Evaluation of clustering performance across 300 configurations combining 30 GLCM window sizes and 10 cluster counts. Each heatmap visualizes the average score across 5-fold cross-validation for a different metric: (a) *F1* score, (b) Dice coefficient, (c) Intersection over Union (IoU). The configuration yielding the best individual performance is marked with a red circle and annotated with its score. While each metric highlights slightly different numerical values, all three exhibit a similar spatial pattern of performance across the parameter space. This consistency reinforces the robustness of the selected configurations and supports the use of *F1* score as a representative criterion for subsequent analysis.

cluster counts and window sizes between 15 and 35—achieve notably high scores across all metrics. These results suggest that certain combinations of texture scale and clustering granularity are especially effective for isolating the salt body. Among these, the configuration with 3 clusters and a window size of 15 was selected based on the highest *F1* score, which served as the criterion for identifying the cluster most predictive of salt. Its corresponding binary segmentation is shown in Figure 4, where salt pixels are displayed in white and all others in black. Although this represents the best individual result, the presence of a broader "island" of high-performing configurations across metrics reinforces the motivation to adopt a consensus-based strategy, integrating multiple informative clusterings into a unified and robust segmentation.

To construct the consensus segmentation map, we first identified the subset of configurations that demonstrated consistently strong performance. Specifically, we selected all parameter combinations whose *F1* score reached at least 95% of the maximum value (0.78) observed across the full grid search. This threshold was chosen to retain only high-performing segmentations while still allowing for some diversity in window size and number of clusters. The resulting subset spans a compact but informative region of the parameter space, capturing multiple textural scales and segmentation granularities that were effective at isolating the salt body. To account for differences in performance among these top configurations, we assigned each one a normalized weight based on its *F1* score relative to the maximum. These weights were later used to modulate each configuration's influence on the final consensus, ensuring that more accurate segmentations contributed proportionally more to the result.

For each of the selected configurations, cluster centroids were obtained by analyzing only the expert-annotated salt regions within the seismic image. Specifically, local texture features were extracted from every pixel inside the salt mask using a window size defined by the current configuration. Each patch was converted into a 4-dimensional feature vector representing contrast, correlation, energy, and homogeneity, computed from GLCMs averaged across eight directional offsets. The resulting set of feature vectors—drawn exclusively from the known salt body—was then clustered using k-means, with the number of clusters k set according to the configuration. To identify which cluster best characterized the salt structure, each cluster's segmentation mask was evaluated against the expert salt label using the *F1* score. The cluster achieving the highest *F1* score was selected as the most representative of salt for that configuration, and its corresponding centroid was retained for subsequent analysis.

Once the salt-representative cluster was identified for each configuration, the full seismic image was evaluated pixel by pixel. At each location, a local patch was extracted using the window size associated with the configuration, and the same 4-dimensional GLCM feature vector averaged across eight directions was computed. This feature vector was then compared to the

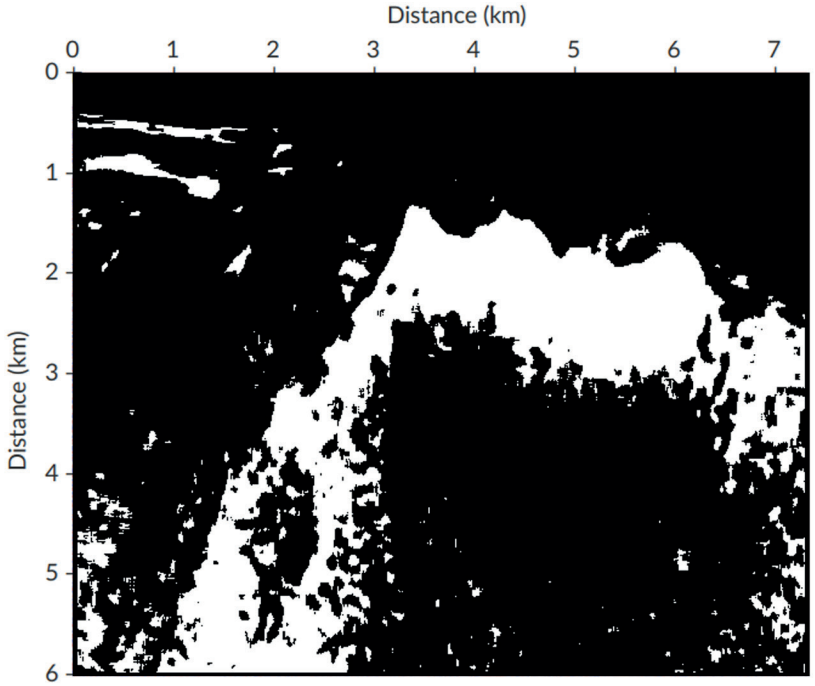

Figure 4. Binary segmentation result from the best-performing configuration (3 clusters, window size 15), as determined by its *F1* score (0.78). Pixels classified as part of the salt body are shown in white, while all other clusters have been merged and displayed in black. This visualization highlights the clustering configuration that achieved the closest match to the expert-defined salt region.

previously computed cluster centroids using Euclidean distance. If the nearest centroid corresponded to the cluster identified as salt-representative, the pixel was considered a match and received a "vote" for belonging to salt. Importantly, each vote was scaled by the configuration's normalized *F1* score, allowing more accurate configurations to contribute more strongly to the final consensus. This process was repeated independently for each configuration, resulting in a weighted accumulation of votes for each pixel across all top-performing models.

As each configuration cast its weighted votes across the image, a cumulative consensus map was constructed by summing the contributions for every pixel. The total vote count at each location reflects the number and strength of high-performing configurations that independently identified that pixel as part of the salt structure. Importantly, no clustering or feature recomputation need to be performed at this stage—each configuration can use the centroids previously derived from the expert-annotated salt regions, ensuring consistency across evaluations. The resulting output is a continuous heatmap, shown in Figure 5, where pixel intensities encode the degree of agreement among top configurations. Brighter areas indicate stronger consensus that a pixel belongs to salt, effectively capturing both spatial detail and segmentation confidence without compromising the original image resolution.

Since no petrophysical measurements from core samples are available for validation, a definitive assessment of the accuracy of the interpretations—manual or automatic—cannot be made. However, the consensus-based result derived from clustering agreement across multiple high-performing configurations exhibits several encouraging features. Notably, the main salt body is strongly and consistently represented in the consensus map, suggesting that salt has a distinctive and quantifiable texture in the GLCM feature space. This reinforces the potential value of the approach as a tool within seismic interpretation software to aid salt body delineation. The top boundary of the salt structure (A) in the consensus map closely aligns with expert interpretations, which is expected given the typically clear seismic reflector associated with this interface. More notably, the base of the salt (B), which is harder to delineate due to weaker and more ambiguous reflections, is captured in the consensus map with a coherent geometry that likely reflects true structural complexity. Whereas manual interpretations tend to favor smoother base geometries, the consensus result suggests a more irregular and geologically plausible shape. Perhaps most compelling is the identification of

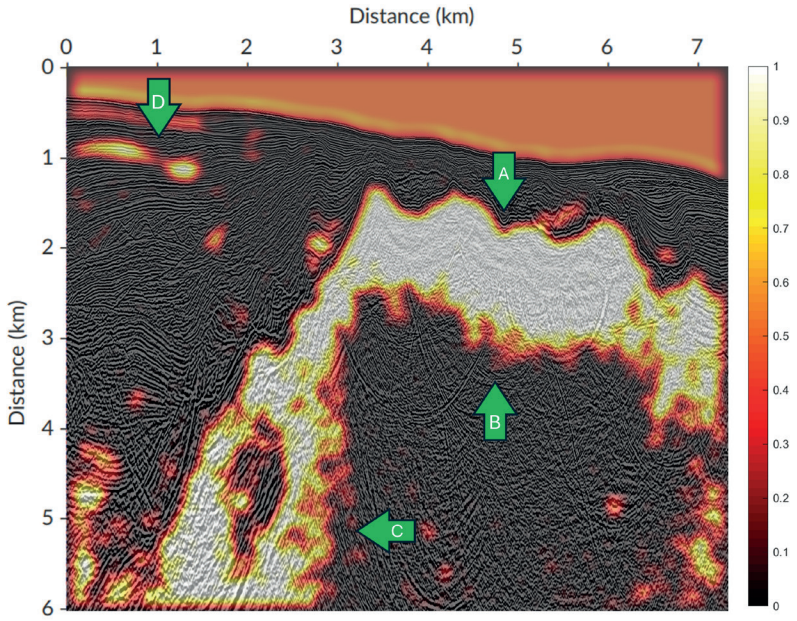

Figure 5. Consensus-based salt delineation heatmap overlaid on the original seismic image. Pixel intensities range from 0 to 1 and represent the normalized confidence level based on the weighted agreement of high-performing clustering configurations. Brighter regions indicate stronger consensus that a pixel belongs to the salt structure. Label A marks the top of the main salt body, showing strong alignment with both expert interpretations. Label B marks the base of the salt structure, where the algorithm produces a well-defined but non-smooth boundary. This irregular contour contrasts with the smoother expert interpretations but may better reflect the expected geological reality, as salt bases are often rough or complex in shape. The result falls between both expert delineations, suggesting the method captures meaningful substructure beyond human smoothing tendencies. Label C highlights a subtle salt-related feature consistently identified by the algorithm but initially missed by human experts, potentially reflecting a salt migration pathway. Label D points to an area of likely over-segmentation, where the method assigns high salt likelihood despite conflicting reflector evidence.

an additional salt-related feature (C) along the left flank of the main body—an area that was overlooked in manual interpretation. The consistency of this feature across multiple configurations in the consensus map implies a genuine textural signature. Further inspection of the surrounding seismic context reveals a system of faults that could plausibly indicate a migration pathway for a salt diapir. At the same time, the consensus result highlights areas of likely over-segmentation or misclassification, such as the top-left corner of the seismic section (D), where the algorithm assigns high salinity confidence to regions unlikely to contain salt based on the surrounding strong reflectors. While these cases reflect limitations of the unsupervised approach, they also provide clear guidance for targeted refinement and expert review.

5. Discussion

The methodology proposed in this study offers a semiautomatic and interpretable approach to salt body delineation in seismic images by combining GLCM texture attributes with unsupervised k-means clustering. Designed to operate under minimal supervision, the method requires only a small set of labeled pixels derived from the intersection of two expert interpretations—thus reducing reliance on large annotated datasets. Its low computational cost and transparent feature set make it especially suitable for early-stage exploration or resource-constrained settings, where complex models may be impractical. Additionally, by aggregating outputs from high-performing configurations using a weighted consensus voting strategy, the method produces a spatially informative heatmap that enhances interpretability. The final segmentation not only aligned closely with expert-labeled salt regions, but also revealed a secondary salt structure overlooked by both interpreters, demonstrating its ability to surface subtle, geologically plausible features that may otherwise be missed.

A known limitation of k-means clustering—particularly when applied to image segmentation—is its lack of spatial awareness. Because the algorithm relies solely on feature similarity and ignores pixel adjacency, it may assign the same label to spatially distant regions, leading to disconnected or geologically implausible segmentations. This behavior is evident in some configurations where salt pixels appear scattered or fragmented. While our use of GLCM-based texture descriptors partially mitigates this issue by encoding some contextual information, the algorithm itself does not enforce spatial coherence. Addressing this limitation represents an important avenue for future work. Possible remedies include incorporating spatial regularization through post-processing techniques, applying models such as Markov Random Fields to promote label continuity, or integrating spatial proximity directly into the clustering algorithm—either through distance-weighted features or spatially informed similarity measures. These enhancements could improve segmentation smoothness while preserving the method's lightweight and interpretable nature.

Another limitation of the proposed approach lies in the inherent sensitivity of *k*-means to both the number of clusters (k) and the random initialization of cluster centroids. Selecting an inappropriate k can lead to undersegmentation, where salt is merged with background sediments, or oversegmentation, where salt regions are fragmented across multiple clusters. Moreover, since *k*-means can converge to local minima, different initializations may yield different results for the same configuration. To mitigate these issues, we performed an exhaustive grid search across a wide range of k values and window sizes and implemented multiple runs per configuration to reduce randomness. Additionally, we adopted a consensus-based voting strategy that aggregates the results of several high-performing configurations, thereby enhancing stability and reducing reliance on any single clustering outcome. Nonetheless, some variability and sensitivity may remain. Future work may explore more sophisticated clustering strategies, such as consensus clustering ensembles or automatic selection of k based on stability metrics, to further improve robustness and reduce the need for manual parameter tuning.

A notable strength of our approach is how it leveraged the contrasting nature of two available expert interpretations—one conservative and one aggressive—to define high-confidence training regions. By intersecting these interpretations, we isolated pixels that were consistently labeled as either salt or non-salt, minimizing the inclusion of ambiguous or debatable areas. This strategy not only reduced potential noise in the training data but also created a buffer zone of uncertainty, which we intentionally left to the unsupervised model to resolve. This division of labor—where clear cases guide clustering while unclear regions are left open to automated exploration—proved effective in both improving model stability and highlighting areas of geological ambiguity. In contexts where extensive annotation is infeasible, this selective use of expert input offers a pragmatic compromise between label reliability and annotation effort, while still enabling the extraction of useful insights from uncertain regions.

Despite its promising results, the present study has several limitations. First, it focuses exclusively on a single two-dimensional seismic section. This inherently restricts volumetric continuity, as interpretations must be made slice by slice without accounting for inter-slice spatial coherence. Consequently, the method's generalizability to full 3D seismic volumes remains untested. Second, the validation was performed using expert interpretations in the absence of petrophysical ground truth such

as well logs or core samples. While expert-labeled data provide a reasonable benchmark, they are inherently subjective and may vary between interpreters. Third, the methodology was applied to a single seismic line featuring salt bodies of moderate depth and geometry. Its performance across a broader range of salt morphologies—such as thin tongues, rugose flanks, or deeply buried isolated bodies—has not yet been evaluated. Future work should therefore explore applications in more diverse geological settings and extend the method to three-dimensional volumes, where spatial context could enhance both clustering fidelity and interpretability. The integration of petrophysical data would also support a more objective assessment of accuracy and geological plausibility. In addition, future research should investigate the influence of seismic data preprocessing—such as stacking, filtering, or gain correction—on the performance and interpretability of GLCM-based clustering, given the sensitivity of texture attributes to signal enhancement or degradation. From a methodological standpoint, however, the proposed approach is not tied to any specific seismic attribute or processing workflow. Because it relies on statistical textural descriptors, which are inherent to virtually all seismic datasets, the method is theoretically applicable across various attribute domains. While the visual expression of patterns may differ depending on the data used (e.g., amplitude, phase, coherence), the ability to extract and cluster meaningful textures remains fundamentally intact.

To build upon the findings of this study, several avenues for future research are recommended. First, extending the methodology to three-dimensional seismic volumes would allow the exploitation of spatial continuity across slices, likely improving both segmentation consistency and geological plausibility. Second, exploring alternative clustering algorithms that explicitly incorporate spatial context—such as superpixel-based methods, region growing techniques, or probabilistic models like Markov Random Fields—could address the spatial disconnects observed with k-means. Third, semi-supervised strategies could be developed where experts iteratively refine consensus segmentations, guiding the model while maintaining interpretability. Fourth, benchmarking the method against state-of-the-art CNNs or Bayesian segmentation approaches would help clarify trade-offs in terms of accuracy, complexity, and explainability. Finally, applying the method to additional datasets from different basins or geological settings would provide insight into its generalizability and robustness across diverse seismic environments.

6. Conclusions

Interpreting salt bodies in seismic data remains one of the most critical and challenging tasks in subsurface exploration, with implications for hydrocarbon prospecting, CO_2 sequestration, and broader geotechnical assessments. The opaque geometry of salt structures and the poor illumination of subsalt regions often lead to ambiguities and inconsistencies in manual interpretation. In response to this challenge, the present study proposed a lightweight, transparent, and semiautomatic methodology that leverages GLCM texture attributes and unsupervised k-means clustering to assist in delineating salt bodies. This approach offers an interpretable alternative to more complex or opaque models, emphasizing accessibility and practical deployment in data-limited environments.

The methodology achieved a maximum *F1* score of 0.78 when benchmarked against high-confidence regions derived from expert interpretations, demonstrating strong alignment with manual delineations. Moreover, the consensus map—constructed by integrating multiple high-performing configurations—proved particularly valuable in ambiguous regions, where it highlighted areas of model agreement and surfaced features overlooked by both human interpreters. This reinforces the method's practical utility as a support tool, rather than a replacement for expert judgment. By offering consistent, reproducible, and interpretable guidance, the approach can accelerate the initial screening of complex salt geometries and reduce subjectivity in early-stage seismic interpretation.

In broader terms, methods like the one proposed here are especially valuable in low-data environments or during the early phases of exploration, where labeled examples are scarce and rapid interpretation is essential. By relying on interpretable features and transparent decision-making processes, this approach aligns with emerging trends in geoscience that emphasize explainability and user trust in AI-driven tools. When integrated into an interactive workflow—where experts refine and validate machine-generated insights—such tools can enhance both the efficiency and the consistency of seismic interpretation. Ultimately, the synergy between computational analysis and human expertise holds the potential to transform the way complex subsurface environments are explored and understood.

7. References

Babakhin, Y., Sanakoyeu, A., & Kitamura, H. (2019). Semi-supervised segmentation of salt bodies in seismic images using an ensemble of convolutional neural networks. [Presentación de paper] *Pattern Recognition: 41st DAGM German Conference, DAGM GCPR 2019, Dortmund, Germany, September 10--13, 2019, Proceedings 41*, 218–231. doi: **https://doi.org/10.1007/978-3-030-33676-9_15**

Berthelot, A., Solberg, A. H. S., & Gelius, L.-J. J. (2013). Texture attributes for detection of salt. *Journal of Applied Geophysics, 88*, 52-69. doi: **https://doi.org/10.1016/j.jappgeo.2012.09.006**

Berthelot, A., Solberg, A. H. S., Morisbak, E., & Gelius, L.-J. (2011). Salt diapirs without well defined boundaries--a feasibility study of semi-automatic detection. *Geophysical Prospecting, 59*(4), 682–696. doi: **https://doi.org/10.1111/j.1365-2478.2011.00950.x**

Cavalin, P., & Oliveira, L. S. (2017). A review of texture classification methods and databases. *2017 30th SIBGRAPI Conference on Graphics, Patterns and Images Tutorials (SIBGRAPI-T)*, 1–8. doi: 10.1109/SIBGRAPI-T.2017.10.

Chopra, S., & Alexeev, V. (2006). Applications of texture attribute analysis to 3D seismic data. *The Leading Edge, 25*(8), 934–940. doi: **https://doi.org/10.1190/1.2335155**

Cohen, I., Coult, N., & Vassiliou, A. A. (2006). Detection and extraction of fault surfaces in 3D seismic data. *Geophysics, 71*(4), P21--P27. doi: **https://doi.org/10.1190/1.2215357**

Couder-Castañeda, C., Orozco-Del-Castillo, M., Padilla-Perez, D., & Medina, I. (2025). A parallel texture-based region-growing algorithm implemented in OpenMP. *Scientific Reports, 15*(1), 5563. doi: **https://doi.org/10.1038/s41598-025-89960-8**

Di, H., Shafiq, M., & AlRegib, G. (2018). Multi-attribute k-means clustering for salt-boundary delineation from three-dimensional seismic data. *Geophysical Journal International, 215*(3), 1999–2007. doi: **https://doi.org/10.1093/gji/ggy376**

Farrokhnia, F., Kahoo, A. R., & Soleimani, M. (2018). Automatic salt dome detection in seismic data by combination of attribute analysis on CRS images and IGU map delineation. *Journal of Applied Geophysics, 159*, 395–407. doi: **https://doi.org/10.1016/j.jappgeo.2018.09.018**

Gao, D. (2004). Texture model regression for effective feature discrimination: Application to seismic facies visualization and interpretation. *Geophysics, 69*(4), 958–967. doi: **https://doi.org/10.1190/1.1778239**

Geng, Z., Hu, Z., Wu, X., Liang, L., & Fomel, S. (2022). Semisupervised salt segmentation using mean teacher. *Interpretation, 10*(3), SE21–SE29. doi: **https://doi.org/10.1190/INT-2021-0191.1**

Hall-Beyer, M. (2000). GLCM texture: a tutorial. *National Council on Geographic Information and Analysis Remote Sensing Core Curriculum, 3*(1), 75. doi: **http://dx.doi.org/10.13140/RG.2.2.12424.21767**

Ji, X., BenHasan, N., Luo, Y., Gashawbeza, E., & Saleh, S. M. (2020). Recognition of salt zones in 3D seismic images using machine learning. *SEG Technical Program Expanded Abstracts, 1626-1630*. doi: **https://doi.org/10.1190/segam2020-3426105.1**

Jones, I. F., & Davison, I. (2014). Seismic imaging in and around salt bodies. Interpretation, 2(4), *SL1-SL20*. doi: **https://doi.org/10.1190/INT-2014-0033.1**

Lomask, J., Biondi, B., & Shragge, J. (2004). Image segmentation for tracking salt boundaries. *SEG Technical Program Expanded Abstracts 2004, October*, 2443–2446. doi: **https://doi.org/10.1190/1.1851242**

Love, P. L., & Simaan, M. (1985). Segmentation of a seismic section using image processing and artificial intelligence techniques. *Pattern Recognition, 18*(6), 409–419. doi: **https://doi.org/10.1016/0031-3203(85)90011-1**

Musgrave, M. J. P. (1960). Reflexion and refraction of plane elastic waves at a plane boundary between aeolotropic media. *Geophysical Journal International, 3*(4), 406–418. doi: **https://doi.org/10.1111/j.1365-246X.1960.tb01714.x**

Ng, I., Kittler, J., & Illingworth, J. (1993). Supervised Segmentation Using a Multiresolution Data Representation. *Signal Processing, 31*(2), 133–163. doi: **https://doi.org/10.1016/0165-1684(93)90062-F**

Orozco-del-Castillo, M. G., Cárdenas-Soto, M., Ortiz-Alemán, C., Couder-Castañeda, C., Urrutia-Fucugauchi, J., & Trujillo-Alcántara, A. (2017). A texture-based region growing algorithm for volume extraction in seismic data. *Geophysical Prospecting*, 65(1), 97–105. doi: **https://doi.org/10.1111/1365-2478.12381**

Orozco-del-Castillo, M. G., Ortiz-Alemán, C., Urrutia-Fucugauchi, J., & Rodríguez-Castellanos, A. (2011). Fuzzy logic and image processing techniques for the interpretation of seismic data. *Journal of Geophysics and Engineering, 8*(2), 185–194. Doi: **https://doi.org/10.1088/1742-2132/8/2/006**

Patterson, J., & Gibson, A. (2017). *Deep learning: A practitioner's approach*. O'Reilly Media, Inc.

Ratcliff, D. W., Gray, S. H., & Whitmore Jr, N. D. (1992). Seismic imaging of salt structures in the Gulf of Mexico. *The Leading Edge, 11*(4), 15–31. doi: **https://doi.org/10.1190/1.1436876**

Steklain, A. F., Ganacim, F., Adames, M. R., Gonçalves, J. L., & Oliveira, D. S. (2022). Structure-oriented filtering in unsupervised multiattribute seismic facies analysis. *The Leading Edge, 41*(6), 411–4171. doi: **https://doi.org/10.1190/tle41060411.1**

Tan, P.-N., Steinbach, M., & Kumar, V. (2016). *Introduction to data mining*. Pearson Education India.

Tavakolizadeh, N., & Bagheri, M. (2022). Multi-attribute Selection for Salt Dome Detection Based on SVM and MLP Machine Learning Techniques. *Natural Resources Research, 31*(1), 353–370. doi: **https://doi.org/10.1007/s11053-021-09973-8**

Wang, Q., Wang, C., Feng, Z., & Ye, J. (2012). Review of K-means clustering algorithm. *Electronic Design Engineering, 20*(7), 21–24.

West, B. P., May, S. R., Eastwood, J. E., & Rossen, C. (2002). Interactive seismic facies classification using textural attributes and neural networks. *The Leading Edge, 21*(10), 1042–1049. doi: **https://doi.org/10.1190/1.1518444**

Xu, Z., Li, K., Yin, R., Fan, Y., & Ma, J. (2025). 3D Saltseg-CL: Unsupervised embedding characterization based multi-task dense prediction method for 3D salt bodies. *Expert Systems with Applications, 267*, 126249. doi: **https://doi.org/10.1016/j.eswa.2024.126249**

Zhang, Z., Zhu, Q., & others. (2012). Fuzzy time series forecasting based on k-means clustering. *Open Journal of Applied Sciences, 2*(04), 100–103. doi: **http://dx.doi.org/10.4236/ojapps.2012.24B024**