WEAKLY-SUPERVISED LABELING OF SEISMIC VOLUMES USING REFERENCE EXEMPLARS

Yazeed Alaudah and Ghassan AlRegib

Center for Energy and Geo Processing (CeGP) at Georgia Tech and KFUPM
School of Electrical and Computer Engineering
Georgia Institute of Technology
{alaudah,alregib}@gatech.edu

ABSTRACT

Localizing seismic structures that can form traps for hydrocarbon reservoirs within large seismic volumes is a very challenging task. Due to the lack of accurately labeled data, we propose a weakly-supervised model for labeling seismic volumes using only a few labeled exemplars. Using six manually-labeled patches, we are able to extract patches that contain instances of similar geophysical structures. Features based on the effective singular values of curvelet coefficients are then used to train a classifier that can label an entire seismic volume with relatively high accuracy. Experiments on reference seismic sections in the Netherlands North Sea seismic dataset result in 73.8% mean pixel accuracy and 75.5% mean class accuracy, with an average labeling time of 5.2 seconds per section. These results are promising considering the nature of seismic images, and the lack of accurate edges between different geological structures.

Index Terms— Computational Seismic Interpretation, Segmentation, Weakly Supervised Labeling, Feature Extraction, Image Classification

1. INTRODUCTION

In the oil and gas industry, seismic surveys are conducted over hundreds of square kilometers to evaluate the prospect of hydrocarbons in specific geological regions. These surveys result in terabytes of data that geophysicists and seismic interpreters have to spend months laboriously analyzing. The process of interpreting post-migration seismic data is very complex. An essential task in this process is identifying and labeling structures that have the possibility of trapping hydrocarbons such as salt domes, and faults. These structures have conventionally been identified and labeled by experienced interpreters, but this process is highly labor intensive and very time consuming.

In recent years, there has been increasing interest in automated or semi-automated tools for the interpretation of seismic data. For example, AlBinHassan and Marfurt [1] proposed using the 2D hough transform to detect fault lines in 2D seismic sections. Wang *et al.* [2] proposed an improved method that removed noisy line features detected by the 2D hough transform, and Jacquemin and Mallet [3] proposed a cascade Hough transform to detect fault surfaces. Also, Admasu *et al.* [4] used log-Gabor filters to denoise along the direction of faults and tracked fault lines in the seismic volume using active contours. In addition, Yan *et al.* [5] proposed using the ant colony algorithm to track faults in seismic volumes, and Wang *et al.* [6] proposed using the idea of motion vectors in video coding to automatically track faults in seismic volumes.

In addition, researchers have proposed different algorithms for detecting and delineating salt dome structures. Aqrawi *et al.* [7] proposed a method for detecting salt domes using a dip-guided 3D Sobel filter. Berthelot *et al.* [8] proposed using three groups of texture features in a supervised Bayesian classification model. Hegazy and AlRegib [9] proposed three texture features for salt dome detection based on the moment of inertia tensor, while Amin and Deriche [10] proposed hybrid edge- and texture-based features. More recently, Shafiq *et al.* [11] proposed using the three-dimensional gradient of textures (3D GoT) for detecting salt dome boundaries while Wang *et al.* [12] proposed a tensor-based subspace learning algorithm for tracking these boundaries in large seismic volumes.

However, many of these algorithms implicitly assume seismic data interpreters would manually segment a region around these geological structures before applying these algorithms. This is because in the process of interpreting large seismic volumes, interpreters first assume a geological model based on the geological history of the region. Then the volume is manually segmented into smaller volumes based on the dominant structure within each sub-volume. Then within each sub-volume, interpreters use various methods (like some of earlier cited methods) to accurately identify and localize key structures. Based on the results, the geological model is modified and the process is repeated until the interpreters converge on a geologically plausible and reasonably accurate model. While the first step of roughly segmenting the large

volume is usually done manually, we show in this work that it can be done automatically, and then various other algorithms specific to certain structures can be applied as a following step. We specifically focus in this work on faults and salt dome structures, given the number of proposed methods devised to delineate them, but naturally this can be extended to other seismic structures as well.

A major challenge in working with seismic data is the lack of accurately-labeled training data. In addition, seismic data lacks the clearly defined boundaries between objects that exist in natural images. For these reasons, we propose using a small number of hand-labeled exemplars to automatically extract and label a large set of images that contains similar geological structures. This process is done using a similarity metric that we propose in [13]. The details of the proposed method is presented in the next section.

2. PROPOSED METHOD

The proposed method can be divided into four steps. First, given a set of exemplar images, a similarity metric is used to extract a large number of images with the same geological structures. This data extraction stage is detailed in section 2.1. Secondly, a set of features is extracted from each image, and are used to train a classifier. This is detailed in section 2.2. Once the classifier is trained, we segment each seismic section in the database, and classify the structures in that section. These steps are detailed in sections 2.3 and 2.4, respectively. A high-level block diagram of the proposed method is shown in Fig. 1.

2.1. Data Extraction

The focus of this paper is on post-migration seismic sections. These are grayscale images devsied by stacking seismic traces which are obtained after time or depth migration of a conventional seismic survey. Sample patches extracted from these images are shown in Fig. 2. We define three structures that are of interest, namely `Chaotic` layers, `Faults`, and `Salt dome`. We also define the `Other` class for patches that don't contain any of the previous three structures. An example from each class is shown in Fig. 2. In our work, we use seismic sections extracted along the crossline direction of the widely used Netherlands North Sea Offshore F3 Block [14]. All sections where normalized to have zero mean and a standard deviation of 1. This is done before extracting the patches to remove contrast and mean variations between different sections. Six exemplar images (two each from `Salt dome` and `Other`, and one each from the rest) were extracted and manually labeled. We then use the texture similarity metric we propose in [13] to extract 500 images that have the most similarity to each exemplar image. Feature vectors extracted from these images are later used to train the classifier.

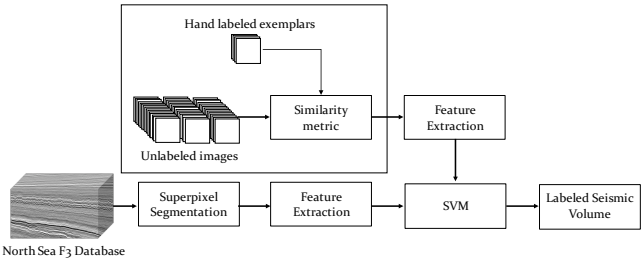

Fig. 1: A high-level block diagram of the proposed method.

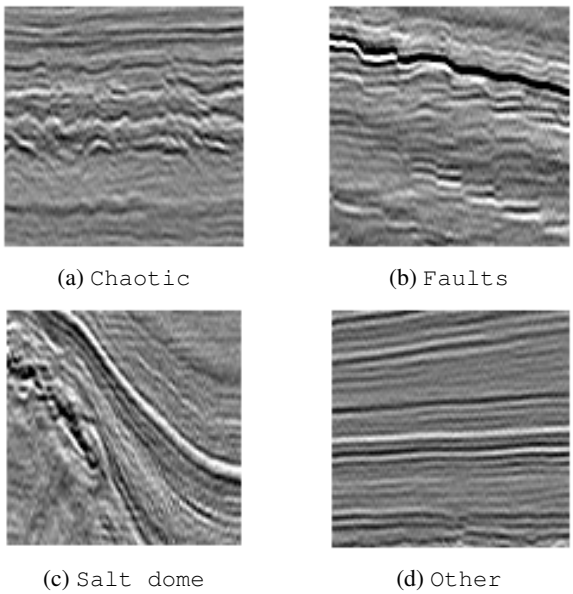

(a) `Chaotic` (b) `Faults`

(c) `Salt dome` (d) `Other`

Fig. 2: Sample images from the four different classes, taken from the Netherlands North Sea F3 Block dataset [14].

2.2. Feature Extraction

In order to train our classifier, and later predict the classes of images that are extracted from the seismic sections, we need to extract relevant features from each image. Given an image X_i, we first take the Hadamard product of this image with a two-dimensional Gaussian kernel of the same size G that is centered at the middle of the image. This is done to emphasize the local spatial correlations in seismic data, and to give more weight to the structures at the center of the image, and less weights to those on the periphery. This can be expressed as

$$\widetilde{X}_i = X_i \odot G, \tag{1}$$

where $\odot$ is the Hadamard product, and G is defined as

$$G[x, y] = e^{\frac{(x-\mu_x)^2+(y-\mu_y)^2}{2\sigma^2}} \tag{2}$$

where μ_x and μ_y are the x- and y- coordinates of the center of X_i respectively. The value of σ is selected such that pixels in the corners of the image have weights of less than 1%.

Next, we represent each image with a feature vector that we propose in [13]. This was shown to be a highly effective representation for texture images in general, and seismic images specifically. This representation is based on taking the singular value decomposition (SVD) of the curvelet [15] coefficients of the image at all scales and orientations, and then zeroing out all singular values larger than the effective rank of every curvelet tile [16]. The final feature vector is formed by concatenating the resulting singular values across all the curvelet orientations and scales. For an image of size 99×99, the resulting feature vector has 500 elements.

Due to the limited number of training images, and the high level of redundancy in the curvelet transform, we use Principal Component Analysis (PCA) to reduce the dimensionality of the feature vectors. The goal of PCA is to find an orthogonal projection of the data into a linear subspace of dimensionality $d < M$ such that it maximizes the variance of the projected data. This new subspace is called the principal subspace. In addition to reducing redundancy, PCA has the added advantage of reducing the number of training samples required, and thus, the computation time consumed by the SVM.

Given the feature vectors, $\mathbf{f}_i$'s, for all the training samples $\{\widetilde{X}_1, \cdots, \widetilde{X}_N\}$, we can construct the matrix $\mathbf{F} \in \mathbb{R}^{M \times N}$, such that:

$$\mathbf{F} = \begin{bmatrix} \mathbf{f}_1 & \mathbf{f}_2 & \cdots & \mathbf{f}_N \end{bmatrix}, \tag{3}$$

where M is the number of features, and N is the number of training samples. We can obtain a matrix $\bar{\mathbf{F}}$ by subtracting each row in $\mathbf{F}$ by the empirical mean of that row. Now, let $\Sigma_{\bar{\mathbf{F}}} \in \mathbb{R}^{M \times M}$ be the covariance matrix of $\bar{\mathbf{F}}$, then using the eigenvalue decomposition of $\Sigma_{\bar{\mathbf{F}}}$, we have:

$$\Sigma_{\bar{\mathbf{F}}} U = U \Lambda \tag{4}$$

where $\Lambda = \mathsf{diag}(\lambda_1, \lambda_2, \cdots, \lambda_M)$ is a diagonal matrix with the eigenvalues of $\Sigma_{\bar{\mathbf{F}}}$ in its diagonal sorted in descending order such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_M$, and $U = [\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_M]$ contains all the eigenvectors of $\Sigma_{\bar{\mathbf{F}}}$. Then we define $\widehat{U} = [\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_d]_{M \times d}$ where $d < M$, and project $\bar{\mathbf{F}}$ onto the principal feature subspace to obtain $\widehat{\mathbf{F}}$:

$$\widehat{\mathbf{F}} = \widehat{U}^T \bar{\mathbf{F}} \tag{5}$$

where $\widehat{\mathbf{F}} \in \mathbb{R}^{d \times N}$, is used to train our classifier along with the labels $\mathbf{y} = [y_1, y_2, \cdots, y_N]$, where y_i corresponds to the class of i^{th} training image X_i.

Later, to classify superpixels within seismic sections, we project the feature vector of each superpixel $\mathbf{g}_i$ onto the same principal feature subspace to get $\widehat{\mathbf{g}}_i \in \mathbb{R}^d$, which is then fed to the classifier to predict a class. In our work, we select $d = 70$ since the first 70 components were found to contain 99.08% of the variance of the data.

2.3. Segmentation

Given a seismic section, we enforce the local spatial correlation of seismic structures using simple linear iterative clustering (SLIC) superpixels [17]. However, since seismic images are grayscale, we cluster the pixels in the $[l, g_x, g_y, x, y]$ space instead of the $[l, a, b, x, y]$ space. Where g_x and g_y refer to the gradient of the image along the x- and y- directions respectively.

2.4. Classification

A Support Vector Machine (SVM) [18] is a powerful binary classification algorithm. It seeks to find the optimal separating hyperplane between two classes by finding the one with the maximum margin. Since we have a multi-class classification problem, we train four hard-margin SVMs with linear kernels using the one-versus-all (OVA) approach. The classifier is trained using $\widehat{\mathbf{F}}$ and its labels $\mathbf{y}$. The resulting training error, and 5-fold cross-validation error were both less than 0.01%. We predict the labels of different superpixels in the seismic section by first extracting a patch of the same size as X_i centered at the centroid of the superpixel, and then classifying its feature vector.

3. EXPERIMENTAL RESULTS

To evaluate our results objectively, four seismic sections were hand labeled with the help of a geophysics expert. We use four evaluation metrics commonly used in the semantic segmentation literature [19]. If we denote the number of pixels from class y_i classified as y_j as n_{ij}, the number of classes as n_c, and the total number of pixels in class y_i as t_i, then we can define:

- pixel accuracy: $\sum_i n_{ii} / \sum_i t_i$
- mean class accuracy: $(1/n_c) \sum_i n_{ii}/t_i$
- mean intersection over union (MIU): $(1/n_c) \sum_i n_{ii}/(t_i + \sum_j n_{ji} - n_{ii})$
- frequency-weighted intersection over union (FWIU): $(\sum_k t_k)^{-1} \sum_i t_i n_{ii}/(t_i + \sum_j n_{ji} - n_{ii})$

Table 1: Results of the evaluation metrics on different seismic sections

Crossline #	Pixel Acc.	Mean Acc.	MIU	FWIU
61	0.687	0.676	0.416	0.580
211	0.748	0.762	0.474	0.643
231	0.744	0.762	0.461	0.641
261	0.774	0.820	0.498	0.670
Mean	**0.738**	**0.755**	**0.462**	**0.634**

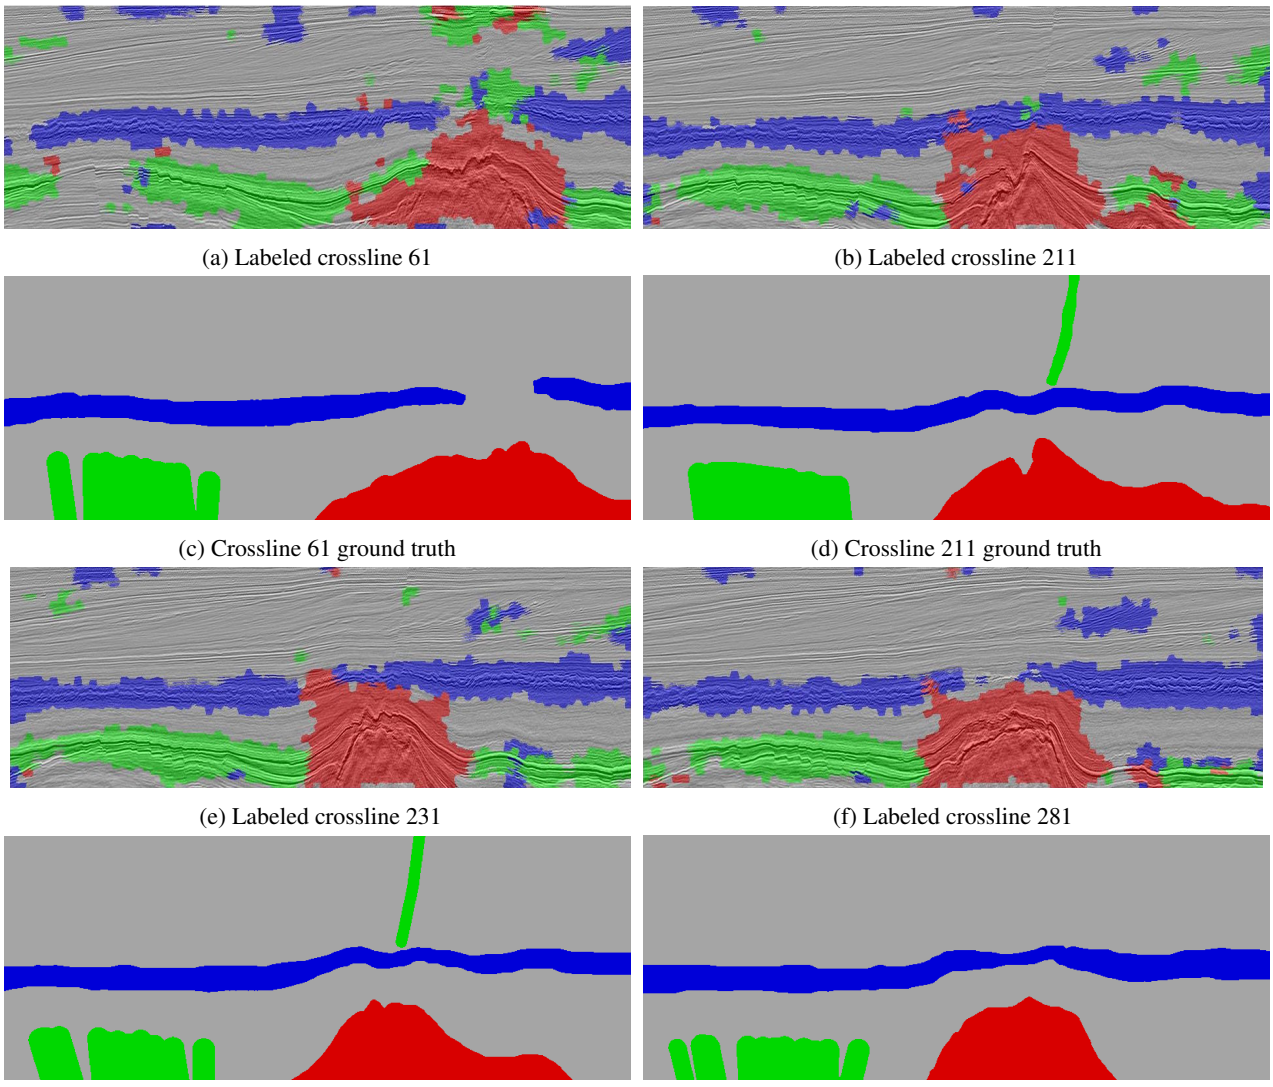

(a) Labeled crossline 61 (b) Labeled crossline 211

(c) Crossline 61 ground truth (d) Crossline 211 ground truth

(e) Labeled crossline 231 (f) Labeled crossline 281

(g) Crossline 231 ground truth (h) Crossline 281 ground truth

Fig. 3: Results of the proposed method on four different seismic sections in the Netherlands North Sea F3 block database [14]. The `Chaotic` class in blue, `Faults` is green, `Salt dome` is red, and `Other` is grey

The objective evaluation results are shown in table 1. Most of the errors in classification are in the `Faults` class. This is especially true in the case where there is only one fault in the middle of the patch. This is case for the large fault in the middle of crosslines 221 and 231. This is partly due to the fact that unlike `Salt dome` features for example, `Faults` are very subtle and small. We are also limited by the single `Faults` exemplar image that was used. The proposed model averages 5.2 seconds to classify a single seismic section of size 390×900 pixels. The simulations were performed using MATLAB 2015b on a PC with the following configuration: Intel core i7 4.0 GHz, with 32GB of RAM, and running 64-bit Windows 10. Considering the very limited labeled data that was used, and the challenging nature of grayscale seismic images, we find these results quite promising.

4. CONCLUSION

In conclusion, we proposed a method for labeling different geophysical structures using only a very small subset of labeled exemplars. A relevant similarity measure is used to extract images that contain similar structures, that are then used as substitutes for labeled data. Features based on the SVD of curvelet coefficients are extracted from these images, and later used to train a support vector machine. Superpixels are used to segment the seismic section, and the SVM is then used to label each superpixel. Objective evaluation shows mean pixel accuracy of 73.8% and mean class accuracy of 75.5%.

5. REFERENCES

[1] Nasher M. AlBinHassan and Kurt Marfurt, "Fault detection using hough transforms," in *SEG Technical Program Expanded Abstracts 2003*, pp. 1719–1721.

[2] Zhen Wang and G. AlRegib, "Fault detection in seismic datasets using hough transform," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 2372–2376.

[3] Pierre Jacquemin and JeanLaurent Mallet, "Automatic faults extraction using double hough transform," in *SEG Technical Program Expanded Abstracts 2005*, pp. 755–758.

[4] Fitsum Admasu, Stefan Back, and Klaus Toennies, "Autotracking of faults on 3d seismic data," *GEOPHYSICS*, vol. 71, no. 6, pp. A49–A53, 2006.

[5] Yan Zhe and Gu Hanming, "An automatic fault tracking approach based on ant colony algorithm," in *SEG Technical Program Expanded Abstracts 2012*, pp. 1–5.

[6] Zhen Wang, Zhiling Long, G. AlRegib, A. Asjad, and M.A. Deriche, "Automatic fault tracking across seismic volumes via tracking vectors," in *Image Processing (ICIP), 2014 IEEE International Conference on*, Oct 2014, pp. 5851–5855.

[7] Ahmed Adnan Aqrawi, Trond Hellem Boe, and Sergio Barros, "Detecting salt domes using a dip guided 3D Sobel seismic attribute," *SEG Technical Program Expanded Abstracts 2011*, pp. 1014–1018, 2011.

[8] Angélique Berthelot, Anne H S Solberg, and Leiv J. Gelius, "Texture attributes for detection of salt," *Journal of Applied Geophysics*, vol. 88, pp. 52–69, 2013.

[9] T Hegazy and G AlRegib, "Texture attributes for detecting salt bodies in seismic data," *2014 SEG Annual Meeting*, vol. 1, no. c, pp. 1455–1459, 2014.

[10] A. Amin and M. Deriche, "A hybrid approach for salt dome detection in 2d and 3d seismic data," in *Image Processing (ICIP), 2015 IEEE International Conference on*, Sept 2015, pp. 2537–2541.

[11] Muhammad A. Shafiq, Zhen Wang*, Asjad Amin, Tamir Hegazy, Mohamed Deriche, and Ghassan AlRegib, "Detection of salt-dome boundary surfaces in migrated seismic volumes using gradient of textures," in *SEG Technical Program Expanded Abstracts 2015*, pp. 1811–1815.

[12] Zhen Wang, Zhiling Long, and G. AlRegib, "Tensor-based subspace learning for tracking salt-dome boundaries," in *Image Processing (ICIP), 2015 IEEE International Conference on*, Sept 2015, pp. 1663–1667.

[13] M. Alfarraj, Y. Alaudah, and G. AlRegib, "Content-adaptive non-parametric texture similarity measure for image retrieval," in *submitted to: Multimedia Signal Processing (MMSP), 2016 IEEE Workshop on*, Sept 2016.

[14] B.V. dGB Earth Sciences, "The Netherlands Offshore, The North Sea, F3 Block - Complete," 1987.

[15] Emmanuel Candes, Laurent Demanet, David Donoho, and Lexing Ying, "Fast discrete curvelet transforms," *Multiscale Modelling and Simulations*, vol. 5, no. 3, pp. 861–899, 2005.

[16] O. Roy and M. Vetterli, "The effective rank: A measure of effective dimensionality," in *Signal Processing Conference, 2007 15th European*, Sept 2007, pp. 606–610.

[17] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 11, pp. 2274–2282, Nov 2012.

[18] V.N. Vapnik, "An overview of statistical learning theory," *IEEE Transactions on Neural Networks*, vol. 10, no. 5, pp. 988–999, Sep 1999.

[19] Jonathan Long, Evan Shelhamer, and Trevor Darrell, "Fully convolutional networks for semantic segmentation," *arXiv*, pp. 3431–3440, 2014.