



Automatic Segmentation of Breakouts in Acoustic Image Logs with Deep Learning

Gabrielle Brandenburg dos Anjos¹, Augusto Ícaro da Cunha¹, Mayara Gomes¹, Nelia Reis¹, Raquel Guilhon¹, Renata Nascimento¹ and Candida Menezes de Jesus², ¹Tecgraf Institute - PUC-Rio, ²Petrobras

Copyright 2023, SBGf - Sociedade Brasileira de Geofísica

This paper was prepared for presentation during the 18th International Congress of the Brazilian Geophysical Society held in Rio de Janeiro, Brazil, 16-19 October 2023.

Contents of this paper were reviewed by the Technical Committee of the 18th International Congress of the Brazilian Geophysical Society and do not necessarily represent any position of the SBGf, its officers or members. Electronic reproduction or storage of any part of this paper for commercial purposes without the written consent of the Brazilian Geophysical Society is prohibited.

Abstract

Identifying breakout failures is fundamental for estimating the borehole's stability and obtaining the direction and magnitude of the maximum horizontal stress in the rock formation. Traditionally, interpreters identify and characterize these collapsed zones manually in acoustic image logs, which can be a very time-consuming and labor-intensive task. Other aspects that make interpretation difficult are the presence of several noisy artifacts and ambiguous pore structures in the image data. With the arrival of deep learning techniques, solutions based on neural networks have become increasingly promising for complex tasks related to object detection and segmentation in images. This work describes a methodology for the segmentation of breakout regions using a convolutional neural network called DC-UNet, and pre and post-processing techniques applied to enhance the quality of the labels. The model achieved an average F-Score of 72.3% and satisfactory qualitative results obtained from a database of 33 acoustic amplitude image logs.

Introduction

In the O&G industry, the analysis of geomechanical stresses is a significant step in reservoir modeling. Drilling a borehole can change the *in-situ* stress of a rock formation and may yield damages or failures in the wellbore. The compressive failures are collapsed zones on the borehole's wall called breakouts. These structures appear on both sides of the well and are 90 degrees azimuth away from the maximum horizontal stress axis. Because of that, their analysis and comprehension are used for building logs of stress orientation and magnitude along the well trajectory (Zoback, 2007).

The acoustic borehole image log is currently the most relevant data for visualizing breakout artifacts. The logging tool emits a sonic pulse to the wall rock and registers the amplitude and transit time of the reflected wave. In these logs, breakouts appear in pairs of vertical and irregular cavities separated by 180 degrees. Interpreting these structures takes a lot of time and effort,

primarily because of log data's massive size and complexity.

Only a few works have been published to detect breakouts automatically. In deep learning, Dias et al. (2020) proposed training a fast-RCNN architecture with synthetic data to detect these features in image logs. The method reports 90% of AUC for synthetic data; however, there's no mention of the model's performance in distinguishing breakouts from other structures in real data. Besides that, the proposed technique only provides the bounding box and not the segmentation of the regions of interest. Another approach suggested by Valentín (2018) is the use of image processing techniques for breakout segmentation in transit time logs. The author profits that breakout areas display a high rate of noise and heterogeneity in the transit time log in order to use a 3x3 standard deviation filter to highlight these regions. The result is then binarized and post-processed with a dilation operation. The foreground candidates are tested to check if they satisfy the geometric and morphological properties of breakouts. The author doesn't mention if this final evaluation is done manually or with any automatic strategy.

In this paper, we propose the supervised training of a convolutional neural network architecture for a pixel-by-pixel classification of breakout areas combined with image processing techniques to enhance the quality of the ground truth and output labels. The database used for training, validation, and testing is composed of 33 acoustic amplitude image logs and their respective ground truth labels.

Method

The methodology developed can be divided into three stages. The first stage of the workflow regards data treatment before feeding the network. The second part deals with the supervised training of the neural network model using the amplitude image log and ground truth mask. The trained model is then used to predict breakout regions from amplitude image logs unknown by the network. The third stage consists of a post-processing step to enhance the predicted regions.

Pre-processing and sampling

The available ground truth labels were generated using a combination of histogram segmentation and manual adjustments made by interpreters, described in Menezes et al. (2016). This method of labeling is very sensitive to noise, therefore the generated foreground regions are inaccurate and disjointed. To attenuate this problem, a

filter was applied to the binary ground truth mask in order to remove all foreground regions with an area smaller than 10 pixels. After that, a dilation morphological operation is applied to expand the boundaries and better delimit the breakout areas, which also helps to join small fragmented regions that should be a single connected component. The result of this processing is illustrated in Figure 1.




Figure 1 - Sample of the annotation mask before (left) and after (right) pre-processing technique. Breakout labels correspond to the green regions.

In order to send the data as input to the CNN model, it is necessary to sample the log data in small squared patches.

A sliding window technique adapted to the cylindrical domain of a wellbore was used to extract these patches. The method is detailed in Anjos et al. (2022, 2023). Besides breakouts, the provided database also contained fractures and vugs annotations. This information was used to discard all the patches which had no examples of breakouts, vugs, or fractures in order to reduce the sparseness of the data.

The chosen model for segmenting the breakout regions from amplitude image samples is the DC-UNet (Lou et al., 2020). This architecture is a variation from the classical U-Net introduced by Ronneberger et al. (2015). Both models were originally proposed for the segmentation of medical images that requires a high level of detail in boundary regions, which is also essential when segmenting cavities in borehole image data. DC-UNet follows the same U-shape structure as its predecessor, with the left side being the encoder stage that extracts features from the image through a succession of convolution operations, and the right side is the decoder that uses transpose convolutions to upsample the features back to the original input shape, as shown in Figure 2.

Convolution blocks of DC-UNet are called Dual Channel Block (DC-Block), and their main purpose is to extract features in multiple scales. The encoder and decoder are connected by concatenating the output of each encoding DC-Block with the input of the decoding DC-Block on the same level in order to combine features of high resolution with low ones, producing a sharper result.

Training the neural network model involves minimizing a loss function that measures the error between the binary output mask from the network and the ground truth annotated mask. The loss function we used was the binary cross-entropy and is defined as follows:

$$cross(y, y') = \Sigma - (y \log(y') + (1 - y) \log(1 - y'))$$

where y is the pixel value predicted by the trained model, and y' is the actual value of the pixel in the ground truth mask.

Post-processing

In our experiments, some predictions resulting from the trained model were fragmented or thinner than expected. To overcome this issue, a post-processing step was included to increase the quality of the results.

The proposed technique consists of a modified version of the image processing amplitude aware region growing algorithm, flood fill (Vadevenne, 2004). The flood fill algorithm is a classical image segmentation method that applies a global predicate to measure similarity over the input image. Starting from a seed and based on this predicate, the algorithm could continue or stop the region from growing. The procedure consists in define globally an upper and lower threshold values t_a, t_b respectively, and if a pixel lies inside the (t_a, t_b) interval it is labeled as belonging to the segmented region. Otherwise, the pixel is labeled as background.

A traditional value to evaluate is the pixel intensity, so, if the pixel intensity $p \in (t_a, t_b)$, the pixel belongs to the segmented region, otherwise, it belongs to the background. Another strategy is to assign a statistical value, like image variance, as a threshold.

Model Architecture

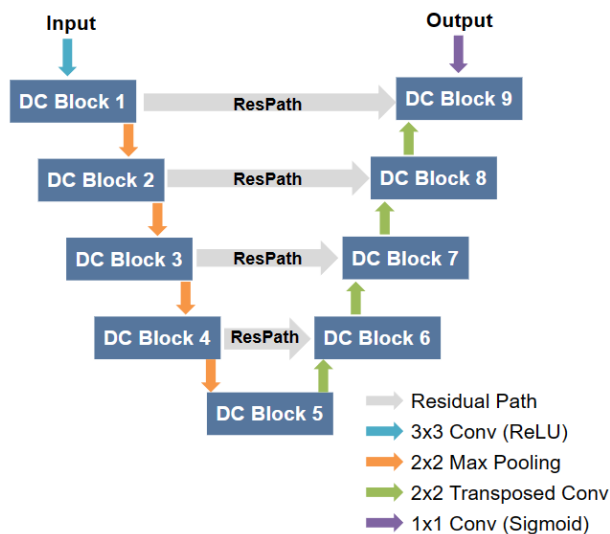


Figure 2 - Architecture of DC-UNet.

The success of this segmentation strategy strongly depends on the correct threshold interval, which must be known a priori. However, finding these limits could be a hard task, particularly for data from acquisitions, like borehole image logs, where the signal can be disturbed by noise. Another issue is that the noise influence along the borehole image logs could vary, so it may be not possible to define a global value to evaluate all pixels.

Proposed strategy

To enhance the automatic segmented structures, the proposed strategy relies on following a local approach instead of a global one. For each detected structure, the initial prediction is taken as a mask. The segmentation mean in terms of intensity value is then computed and the pixel within the labeled ones with the intensity closest to the segmentation mean is taken as seed.

In order to minimize interference from other structures, the image area is reduced to the Region of Interest (ROI). From the initial prediction, a bounding box is extracted as an initial ROI estimation for all predicted structures bigger than 10 pixels. Sequentially, the bounding box is enlarged by 20 pixels at each side, resulting in the final ROI as shown in Figure 3. Once all ROIs are extracted, they are mapped to the image log as a mask to restrict the flood fill propagation.




Figure 3 - ROI acquisition. From left to right: Structures predicted by the trained model, small structures filtering, bounding box extraction, bounding box enlargement.

The following step is to perform a 4-connected flood fill in the amplitude image log. To prepare the data, a Gaussian Blur with a 3x3 kernel size is applied to attenuate the structure's boundary noise. The Otsu's (Otsu, 1979) method is applied in the sequence to binarize the image and then the Canny edge detection (Canny, 1986) is performed to extract the structure's contour. If the resulting contours have gaps, a neighborhood connectivity strategy is applied to connect the extreme points. This pipeline can be seen in Figure 4.




Figure 4 - Structure contours extraction. From left to right: Structure in amplitude image well log, contours extraction, extreme points detection, connected contours.

Once the contour is fully connected, the flood fill is performed starting from the selected seed. In this proposed approach, a pixel p belongs to the final structure S if:

$$M - m * \sigma < I(p) < M + m * \sigma,$$

where $I(p)$ is the intensity of the pixel p , M is the prediction mean, σ is the standard deviation and m is a confidence value set as 2.5 for all experiments.

To guarantee the propagation never crosses the structure contour, all pixel's intensity value within the contour was set to `max_float`. Figure 5 illustrates the post-processing overview for a single structure.




Figure 5 - Post-processing overview. From top left: Structure predicted by the trained model, ROI acquisition, structure selection in amplitude image well log, contour extraction, contour as flood fill constraint, final structure segmentation.

It should be noted that the Otsu algorithm may not be able to always produce the desired binary image version especially when there is a lower contrast variation between the structure and the background. However, in the performed experiments, if the Otsu is applied to small ROIs the results were satisfactory in 99% of the cases.

Results and Discussion

For the quantitative evaluation of the model we used four conventional metrics for classification problems: Recall, Precision, F1-Score, and IoU. In probabilistic terms, Recall is the probability of the model classifying a pixel sample as a breakout given that it is, in fact, a breakout. Precision measures the probability of the output prediction of a pixel sample being correct given that it was classified as a breakout by the model (Goutte et al., 2005). F1-Score is the harmonic mean between Recall and Precision. IoU (Intersection over Union) is the overlapped area of ground truth and predicted mask, divided by the union of these areas.

K-Fold cross-validation is a well-known technique for evaluating a classification model's generalization capacity. It consists of partitioning the dataset into mutually exclusive k subsets, called folds, and running the model k times, using a different subset for training each time.

For our purpose, we created three folds of train, validation, and test, using data from 33 wellbores, described in the table below:

Partition	Fold 1	Fold 2	Fold 3
Test	6 wells (20.9%)	4 wells (21.2%)	7 wells (18.2%)
Validation	7 wells (20.0%)	7 wells (19.9%)	4 wells (21.3%)
Training	20 wells (59.1%)	22 wells (58.9%)	22 wells (60.5%)

Table 1 - Dataset fold description. The number of wells used in each partition and the percentual value of total data volume.

The scores obtained for each evaluated fold and the total score average are shown in Table 2.

The model achieved a good overall performance, with a 72.3% of F1-Score of average. Precision and Recall are well-balanced for Fold 1 and Fold 3, while Fold 2 has the highest Precision and lowest Recall score, caused by the rise of false negatives predictions. Fold 2 has also the biggest test and smaller training set, which correlates with the metrics results. Since the test dataset also contained input images of fractures and vug pores, the results show that the model has a satisfactory capability to distinguish breakouts from other pore structures.

Metrics	Fold 1	Fold 2	Fold 3	Avg
Precision	75.75%	82.0%	79.6%	79.1%
Recall	70.97%	54.9%	76.0%	67.3%
F1-Score	73.28%	65.8%	77.8%	72.3%
IoU	57.83%	49.0%	63.6%	56.8%

Table 2 - DC-UNet scores for each evaluated fold and average.

Some results are shown in Figures 6 and 7 below. Figure 6 shows an example of a successful segmentation of a well-behaved breakout section in an amplitude log. In this result, the predicted structures were very similar to the ground truth. Figure 7 illustrates a section where the model was not able to predict the breakout in the upper region with high levels of noise. However, the breakout in the cleaner zone was successfully predicted, even though it has a less conventional and non-elongated shape.

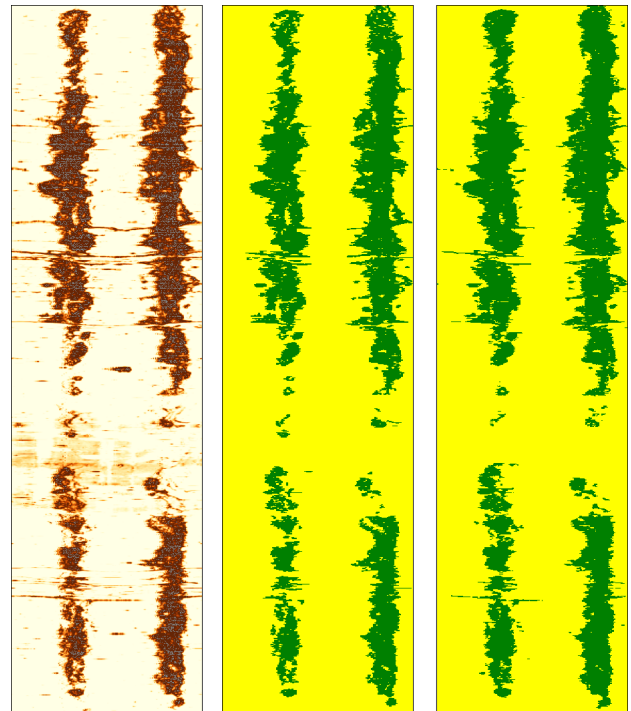


Figure 6 - Example of highly accurate segmentation of a breakout. From left to right: Amplitude image input; ground truth; the output of the proposed method.

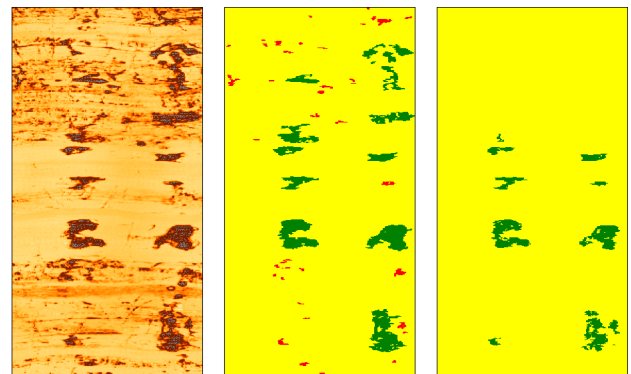


Figure 7 - Example of mid-accurate segmentation of a breakout in a noisy region. From left to right: Amplitude image input; ground truth; the output of the proposed method. Labels in red are related to vug pores.

Conclusions

This work presented a method for the automatic segmentation of breakout areas in acoustic borehole image logs. The methodology consists of supervised training of a convolutional neural network called DC-UNet, followed by a post-processing treatment on the model's prediction output using a flood fill-based technique contained within ROIs.

To validate and evaluate the method we used a dataset of 33 amplitude image logs and their ground truth masks, containing breakout, fractures and vugs examples. The metrics used for quantitative evaluation were Precision, Recall, F1-Score and IoU. The model achieved satisfactory results with an average F1-Score of 72.3% and IoU of 56.8%, and qualitative results that corroborated with the numerical ones.

For future works, we suggest the exploration of different approaches for extracting the patches from the well log, using random cropping for the cylindrical domain. It is also interesting to study how to introduce the transit time log data as a second input channel of the neural network. Another addition would be the use of data augmentation techniques to enrich the dataset, as it seems that the metrics are highly affected by the number of available data for training. The main challenge of the discussed problem is the scarce amount of annotated data, so it is strongly suggested the exploration of unsupervised or semi supervised based solutions.

Acknowledgments

The authors would like to thank Petrobras for the financial support.

References

- ANJOS, G. B.; GATTASS, M.; CUNHA, A. I.; MENEZES, C.; SANTOS, L. F.; GOMES, M.; REIS, N.; GUILHON, R.; NASCIMENTO, R. 2022. Automatic segmentation of breakouts in acoustic borehole image logs using convolutional neural networks. In: Proceedings of the XLIII Ibero-Latin-American Congress on Computational Methods in Engineering, Nov. 2022.
- ANJOS, G. B. 2023. Segmentação Automática de Breakouts em Perfis de Imagem com Aprendizado Profundo. Master thesis, 2023.
- CANNY, J. 1986. A Computational Approach To Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.
- DIAS, L. O. D.; BOM, C. R.; FARIA, E. L.; VALENTÍN, M. B.; CORREIA, M. D.; ALBUQUERQUE, M. P.; COELHO, J. M. 2020. Automatic detection of fractures and breakouts patterns in acoustic borehole image logs using fast-region convolutional neural networks. *Journal of Petroleum Science and Engineering*, v. 191, 2020.
- GOUTTE, C.; GAUSSIER, E. 2005. A probabilistic interpretation of precision, recall and f-score, with implication for evaluation. *Lecture Notes in Computer Science*, p. 345–359, 2005.
- LOU, A.; GUAN, S.; LOEW, M. 2020. Dc-unet: Rethinking the u-net architecture with dual channel efficient cnn for medical images segmentation, 2020.
- MENEZES, C.; COMPAN, A.; SURMAS, R. 2016. Permeability estimation using ultrasonic borehole image logs in dual-porosity carbonate reservoirs. *PETROPHYSICS*, 57:620–637, 12 2016.
- OTSU, N. 1979. A threshold selection method from gray-level histograms. *IEEE Trans. Sys. Man. Cyber.* 9 (1): 62–66, 1979.
- VALENTIN, M. B. 2018. Deep learning methods on geological reservoir borehole log images and applications. Master thesis, 2018.
- VANDEVENNE, L. 2004. Lode's computer graphics tutorial Food Fill, 2004.
- RONNEBERGER, O.; FISCHER, P.; BROX, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, p. 234–241, Cham, 2015.
- ZOBACK, M. D. 2007. Compressive and tensile failures in vertical wells. *Reservoir Geomechanics: Cambridge University Press*, 2007. p. 167–205, 2007.