Automatic Fracture Detection and Characterization in Borehole Images Using Deep Learning-Based Semantic Segmentation

Andrei Baraian[1], Vili Kellokumpu[1], Räty Tomi[1] and Leena Kallio[2]
[1]*VTT Technical Research Centre of Finland, Kaitoväylä 1, Oulu, Finland*
[2]*Astrock Oy, Ahventie 4, Espoo, Finland*

Keywords: Semantic Segmentation, Borehole Analysis, DeepLab, Deep Neural Networks.

Abstract: Fracture analysis represents one of the key investigations that needs to be carried in borehole logs. Identifying fractures, as well as other similar features (like breakouts or foliations) is essential for characterizing the reservoir where the drilling took place. However, identifying and characterizing the fractures from borehole images is a very time and resource consuming task, that require extensive knowledge from geological experts. For this reason, developing semi-automated or automated tools would facilitate and increase the productivity of fracture analysis, since even for one reservoir, experts need to analyze and interpret hundreds of meters of borehole images. This paper presents a deep learning based approach for application of automatic fracture detection and characterization in borehole images, relying on state-of-the-art convolutional neural network for accurate semantic segmentation of fractures. Target images consists of color borehole images, as opposed to acoustic or drill-core images, and uses real world data, both for training the deep learning model and testing the whole system. The system is evaluated by using multiple metrics and the final outputs of the system are the parameters of the sinusoids that define the predicted fractures.

1 INTRODUCTION

An important part of borehole geophysics is the analysis and interpretation of logged measurements for determining the physical properties of wells or test holes. Specific events, such as natural fractures, beddings or foliations approximate the underlying structure of the soil and offer valuable information, depending on the scope of the borehole. For example, micro-fractures in shale allow for identifying the main route for hydrocarbon, which is a clear indicator of oil or gas reservoirs (Carey et al., 2015). Image well logging is one of the most used technique for borehole analysis in our days, and can be implemented in several ways, obtaining electrical, acoustic, or optical images of the borehole. Another popular method of analysing fractures is extracting the drill-core and imaging it at ground-level, using specialized hardware for rolling the core. Although the obtained images are quite similar in terms of content, imaging in borehole is more difficult due to limited illumination configurations, poor imaging conditions in the presence of mud or sediments and prone to image artefacts from camera motion. In our work, the primary focus is on RGB borehole images captured by cameras, but it is possible to transfer our algorithm

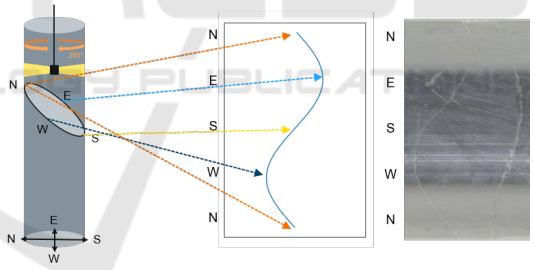

Figure 1: Borehole imaging concept. Fractures in the 3D world are represented as planes cutting through a section of the borehole, generating sinusoid-shaped curves in the unrolled 2D images.

to acoustic or electrical images (Wedge et al., 2015), since the images have a high degree of similarity, even though acquisition methods are different.

To measure and analyse the borehole, a probe is lowered into the borehole and images are continuously captured with a wide FOV (Field-of-View) camera. Images can then be unrolled and stitched together to retrieve a 2D image representation of the inner surface of a borehole. Due to the cylindrical shape of the borehole, planar features such as fractures and beddings will be represented as sinusoidal curves in the unrolled 2D image, as depicted in Fig. 1. Since one borehole can have several hundreds of meters in

Baraian, A., Kellokumpu, V., Tomi, R. and Kallio, L.
Automatic Fracture Detection and Characterization in Borehole Images Using Deep Learning-Based Semantic Segmentation.
DOI: 10.5220/0011673100003417
In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2023) - Volume 4: VISAPP*, pages 856-863
ISBN: 978-989-758-634-7; ISSN: 2184-4321

depth, manually analyzing all the images easily becomes a laborious and time-consuming task for geologists. Our proposed solution is an automated system, capable of accurately and robustly detecting fractures from borehole images and characterizing them at the same time in borehole space parameters. Our solution is developed with data from an existing real world application where fractures are manually annotated with sinusoid parameters. The system is built in such a way to allow experts to set specific confidence score thresholds if more detections are needed.

There are two main approaches for solving fracture detection: traditional image processing techniques combined with pattern recognition, and newer approaches based on Deep Neural Networks (DNNs) (Chen et al., 2018). A generic pipeline for fracture detection consists of two stages: segmenting fracture pixels from the background pixels, and then approximating the models of the curves describing the fractures. We have identified the bottleneck of this system as being the segmentation stage. High quality fracture masks make the fitting easy, however in practice it is exceptionally hard to perfectly segment fractures due to their variability in width, texture, contrast and shape, which forces us to design more complex fitting algorithms. Using traditional image processing techniques, it is very hard and in some cases impossible to derive the perfect combination of filtering operations capable of robustly segmenting the fractures. Even methods relying on more advanced algorithms, such as K-means or kNN clustering are having problems with deep boreholes where color intensity distributions vary a lot. Nonetheless, recent advances in deep learning architectures, especially the ones developed for the task of semantic segmentation, are proving to achieve remarkable accuracy, robustness and performance in scene understanding and classification, being successfully applied in domains such as autonomous driving, optical quality inspection, robot navigation, etc. (Li et al., 2021; Tabernik et al., 2019). This serves as inspiration in our work for choosing DeepLabv3+ (Chen et al., 2018) for the semantic segmentation of fractures.

2 RELATED WORK

Detailed characterisation of borehole features, such as fractures, beddings, foliations, etc. can greatly benefit from the automatization of borehole analysis. Although the research in this domain is still in its beginnings, there are several works that have further pushed the state-of-the-art in borehole analysis and interpretation.

Although the focus in this paper is on optical images captured with color cameras, techniques developed for acoustic or resistivity images can be applied as well. One of the earlier methods proposed by Thapa et al. (Thapa et al., 1997) was relying on computing an edge map and searching for sinusoid candidates by applying the Hough transform. Along the same direction with subtle differences, (Glossop et al., 1999), (Zhang and Xiao, 2009), and (Hall et al., 1996) have applied Hough transform for sinusoid fitting. One of the major drawbacks of HT-based implementations is the high computational cost since it is a brute-force method. In (Moran et al., 2020), the authors aim to overcome this limitation by performing a heuristic search instead of an exhaustive one and leverage the Iterated Local Search (ILS) algorithm. Other geometric transformations have also been applied, such as (van Ginkel et al., 2003), where a general Radon transform is applied. These methods are identifying features that represent fractures, and then by using the geometric transforms, votes are accumulated from each feature and the parameter values with highest support are selected to be the sinusoids. An interesting approach by D. Wedge et al.(Wedge et al., 2015) first highlights regions in the hole most suitable for automated fracture analysis based on an image complexity (quality) measure, arguing that images with a high degree of noise may be ignored even in the manual labeling process. Their method is implemented both for acoustic televiewer and optical televiewer, with the only difference of changing the complexity function.

Recently, methods which leverage the latest advancements in machine learning (ML) are being successfully applied for fracture detection. In (Anatoli Quintanilla Cruz et al., 2017), although the Hough transform is still applied for getting all possible candidates, there is an additional step which consists of the candidates going through a convolutional neural network for validation or exclusion. Dias et al. (Dias et al., 2020) pave the way towards end-to-end systems for fracture and breakouts analysis, by utilizing Fast-RCNN (Girshick, 2015) network to identify fractures and breakouts. To overcome the challenge of limited data, they generate synthetic data of fractures and breakouts and train the model using this data. Although they reach a 98% area under ROC curve in the case of testing on the simulated data, when tested on real images, the system's performance drops significantly. Moreover, only identifying fractures is not enough for a complete autonomous system for borehole analysis, as fractures need to be characterized in terms of borehole parameters such as depth, dip and dip direction. Developed simultaneously with our

work, some ideas from (Alzubaidi et al., 2022) are quite similar to ours, although there are key differences both in the system design and implementation. In terms of imaging, they are using unwrapped core images whereas we use borehole images. To combine the fracture segmentation and instance detection, they are using Mask R-CNN (He et al., 2017) network for instance segmentation of fractures. Although the approach works for the general case and manages to jointly segment and separate each fracture instance, some extreme cases where fractures have a high amplitude might cause problems. After obtaining the mask of the fracture, the next step consists in skeletonizing the fracture, so that the fracture has a width of just one pixel, allowing for better and faster curve fitting. We employ the same idea, however, we perform joint separation where multiple fractures might intersect, enabling us to fit multiple configurations of sinusoids, accounting for sinusoids with high amplitudes as well.

3 DATA

The development and testing of this work was carried on three separate datasets, each depicting images from three different boreholes. The data and annotations are both from a real world application. The resolution of the images was 720x2000 pixels and 1 meter of the borehole was comprised in one image. Each borehole had around 200-300 meters in depth, resulting in 200-300 images per dataset. The annotation of fractures was provided in a separate file for each dataset, having the fractures specified as sinusoids in borehole parameters, hence a conversion to image space was needed in order to generate ground-truth masks. For each dataset, the data was split in training, validation and testing splits with a ratio of 60-20-20 % out of all the images present in a dataset. Consequently, to characterize the predicted sinusoids, it is necessary to perform the inverse operation, from image space to borehole parameters. Because fractures are of different widths, empirically a width of 10 pixels was chosen as the default width of the ground-truth. In practice, the observed fractures are not exactly sinusoids and in borehole images the sinusoid is often only partially visible. Accurate pixel level annotations would be ideal for training the semantic segmentation but annotating data in pixel level is not feasible in the target application. Therefore we focus our study to the use of parametric sinusoid labels produced in the current manual workflow of the application. We think this is a decent tradeoff, as manually labelling and correcting fractures would be infeasible in real world use. The problem is formulated as a two-class semantic segmentation problem, with the classes of fracture and background. Fig. 2 shows images from the three boreholes from where one can see the big variation in terms of texture. Other challenges in the image data include low level artifacts due to borehole image computation from camera images and wide vertical line segments where water blocks the borehole texture as seen in Fig 2.

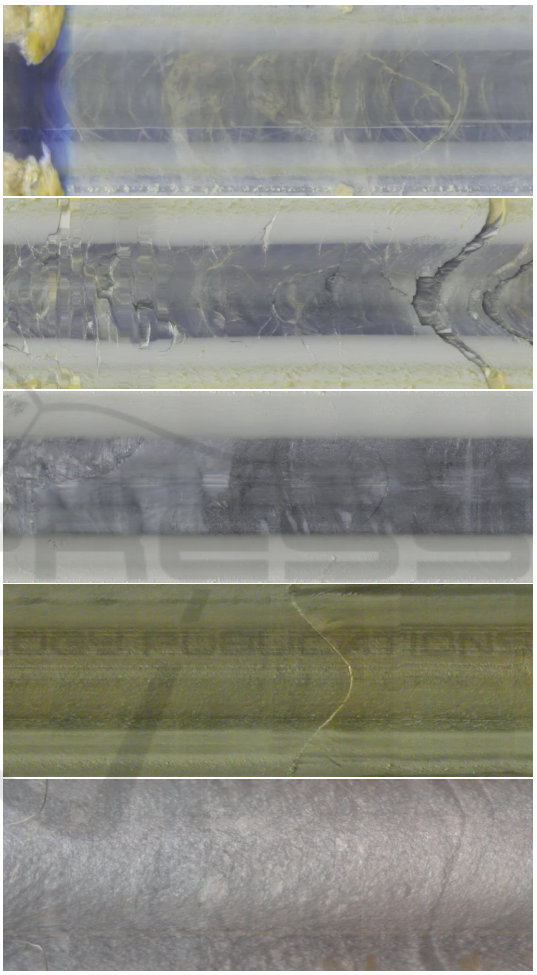

Figure 2: Example images from boreholes.

4 METHOD

Our proposed method relies on two important stages: fracture detection using convolutional neural networks for semantic segmentation and sinusoid fitting on detected fractures. The developed system is aimed at not only detection of fractures, but also accurate identification and characterization of each fracture.

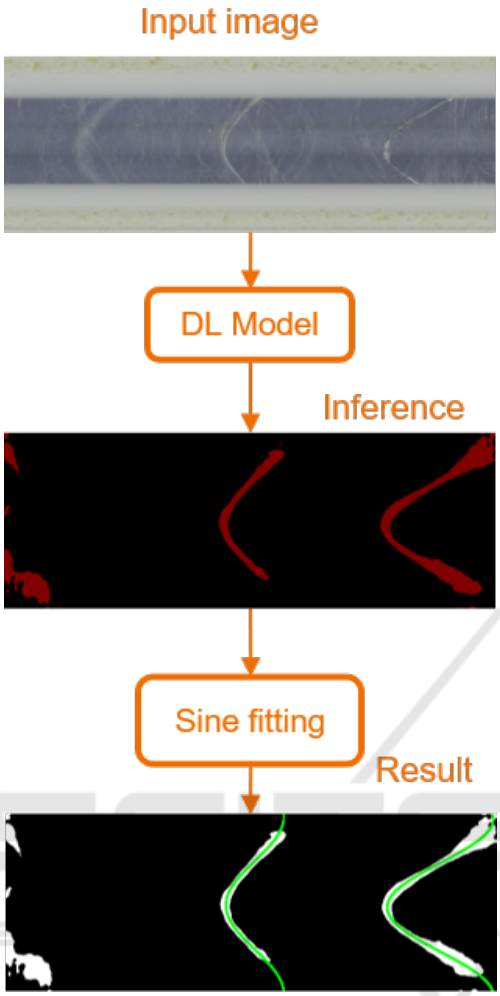

Figure 3: Workflow overview. Semantic segmentation and sinusoid fitting are the two main components of the system.

4.1 Semantic Segmentation

DeepLab is a state-of-the-art Deep Convolutional Neural Network (DCNN) framework developed by Chen et al. (Chen et al., 2018) for the task of semantic segmentation in images. The first version of the framework (Chen et al., 2015) proposed the usage of atrous convolutions and fully connected Conditional Random Fields (CRF) for overcoming the low accuracy in accurate object segmentation. By combining the final DCNN layer with the fully connected CRF, the model is able to bypass the invariance properties of DCNNs, allowing for increased accuracy in segmentation. Last iteration of the framework (Chen et al., 2018) introduces a decoder module for the purpose of refining the segmentation results, especially along the object boundaries. The work is inspired by combining the advantages from the spatial pyramid pooling module and the encoder-decoder structure, both heavily used in the semantic segmentation task. Moreover, the authors adapt the Aligned Xception model (Chollet, 2017) for the semantic segmentation task and apply depthwise separable convolution to both the ASPP and decoder module.

For the training procedure, we have used the pre-trained Xception model from the authors' repository and fine-tuned on our training data, for each dataset. The Xception backbone is initially pre-trained on ImageNet-1k dataset (Russakovsky et al., 2015) and then the DeepLabv3+ model is further pre-trained on MS-COCO dataset (Lin et al., 2014) and on PASCAL VOC 2012 (Everingham et al.,), for which the evaluation is also performed. Because we are more interested in achieving higher accuracy rather than real-time performance, we drop the experimentation on MobileNetV2 (Sandler et al., 2018) and MobileNetV3 (Howard et al., 2019). We run the model with an output stride of 16, hence the atrous rates are [6, 12, 18]. Due to the large image size, we are training with a batch size of 1 image. For fine-tuning, all trained weights are used, except the logits, since the number of classes is different than in PASCAL dataset.

Because the density of fractures in a borehole is quite low, as well as the total number of fractures, there is a high class imbalance between the fracture class and the background class. To alleviate this problem, class weights have been calculated using 1 and applied in the training procedure.

$$W_j = N_s/(N_{cls} * N_{sj}), \quad (1)$$

where N_s is the total number of samples, N_{cls} is the number of classes and N_{sj} represents the number of samples from class j. The calculated weights are then used in the model's loss function.

4.2 Sinusoid Fitting

After getting the masks of possible fractures, the next goal is to fit sinusoid curves to these masks, so that fractures are characterised by curve parameters. The main steps for fracture identification consist in the following steps: skeletonization of the predicted fracture masks, separation of fracture intersection points (joints), connected components labelling of independent fracture segments and brute-force sinusoid fitting of multiple independent skeletons. The output of the semantic segmentation inference consists of a masked image of the identified pixels representing a fracture. However, due to imperfections in segmentation, direct fitting on the predicted masks performs poorly, both in accuracy and computation time. By skeletonizing the fracture masks, their width is de-

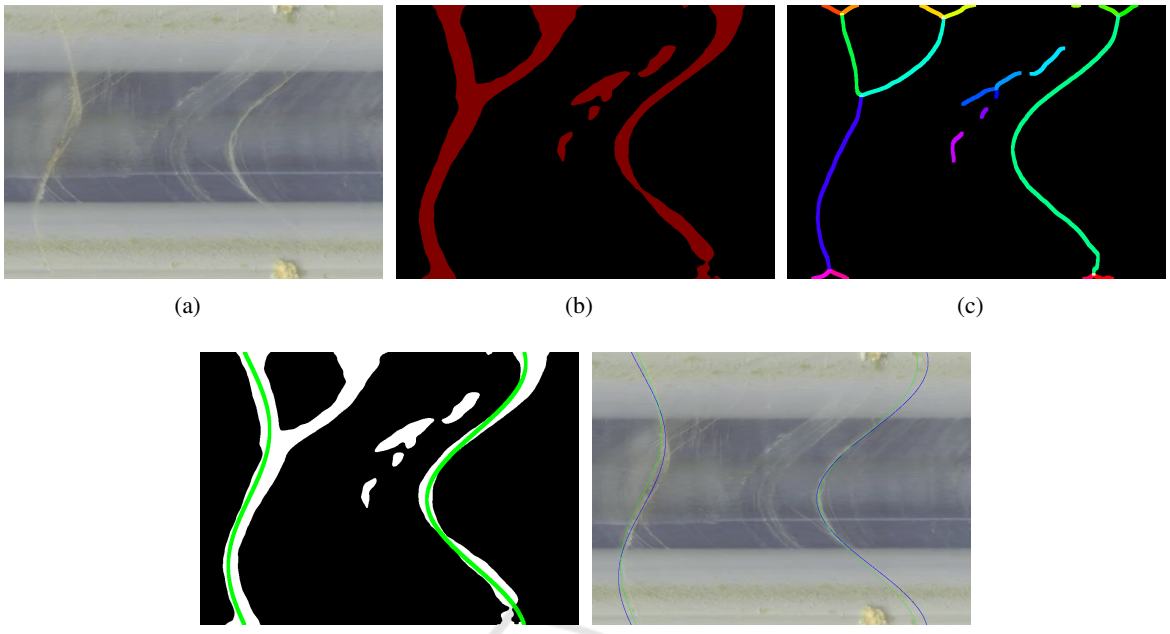

Figure 4: Processing pipeline. (a) Original image. (b) Semantic segmentation of fractures. (c) Fracture mask is skeletonized, joints (intersections of fractures) are removed and independent fracture segments are labelled. (d) Fitted sinusoids on the mask. (e) Sinusoids overlaied on the original image. Blue - GT sinusoids, Green - Detected sinusoids.

creased, however the shape information is still preserved and the number of sample points for the sinusoid fitting is significantly reduced. The skeleton will have a width of 1 pixel. Furthermore, joints in the skeleton are identified and removed, using a simple heuristic based on the number of neighbours. Joint points are detected if there are at least 4 neigbours in an 8-neighborhood region. After the joint points are removed from the skeletons, a connected component labeling algorithm is performed, and each individual skeleton is segmented. Smallest skeleton-fractures are removed and then a brute-force approach is used to create pairs of skeleton-fractures. The fitting is performed on the fracture pair, as well as considering independently each one of the fractures, to account for the case where only one skeleton-fracture defines the whole fracture. To limit the redundant cases, we match skeleton-fractures with fractures which are not on the same horizontal level. The fitting problem is modeled as a simple curve fitting problem in the image space. We assume the case of a vertical borehole, hence we can approximate the shape of the fractures with a sinusoid curve, having the following equation:

$$y = A * sin(B * (x + H)) + V, \tag{2}$$

where A is the amplitude of the sinusoid, $\frac{2*\pi}{B}$ is the period, although for our case we know that the fractures are sinusoids having only one period, so B=1 for our case, H is the phase shift and V is the (depth) vertical shift. The fitting is performed by using SciPy's non-linear least squares optimization library (Virtanen et al., 2020). We are using the Levenberg-Marquardt algorithm without any bounds added. As in any unconstrained optimization problem, we need to estimate an initial guess. For the amplitude and phase shift, we initialise them with a value of 1 and for the vertical shift, we calculate the histogram of the mask on columns and pick the location of the maximum value as the first guess for the vertical shift.

Once the fitting is done and all plausible sinusoids have been computed, we run a Non-Maximal Suppression (NMS) on the detected sinusoids to discard overlapping instances. The detected sinusoid has only 1 pixel width, hence it's width is increased by a fixed amount to account for robustness and overlayed on the initial masks generated by the segmentation, from where a confidence score is derived, based on the overlapping between the detected sinusoid and the generated mask. Then, in a greedy way, only the best sinusoids are retained.

Ultimately, the detected sinusoids parameters are converted from the image-space into the borehole space, in terms of dip angle, dip direction and fracture depth. To calculate the dip angle, the following equation is used:

$$\omega = \arctan(\frac{D}{D_2 - D_1}), \tag{3}$$

where ω is the dip angle, D represents the diameter of the borehole, D_2 is the depth of the upper angle and D_1 is the depth of the lower angle. The dip direction is easily computed by identifying the smaller angle of the amplitudes and converting it from pixels to degrees.

5 EXPERIMENTAL RESULTS

To assess the performance of our system, a multitude of metrics have been applied, to ensure reliable measurements. The performance of the semantic segmentation model can be established by using the mIoU metric, however this is heavily dependent on the ground-truth masks, which were obtained by plotting the sinusoids with a fixed width.

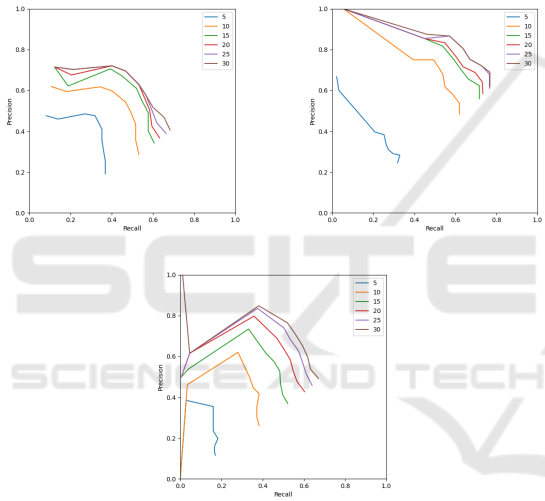

Figure 5: Precision and recall for the 3 datasets.

In practice, fractures are not perfectly following a sinusoidal shape, but there can be slight deviations. Moreover, the width of real fractures varies and establishing such an accurate ground-truth is a laborious work, which is beyond the scope of this work. Hence, we used the mIoU just a mere indication of how to choose some training parameters, such as label weights, atrous rates or batch size. To get an insight on how key parameters are influencing the results, the Precision-Recall curve is an excellent instrument that allows the user to select the optimal tradeoff in terms of sensitivity and specificity.

To derive the PR curve, we first need to classify all detections and ground-truth instances into true positives (TP), false positives (FP). Additionally, each detection needs a confidence score, so that we can choose the optimal tradeoff. We define the confidence score as the number of inliers obtained from

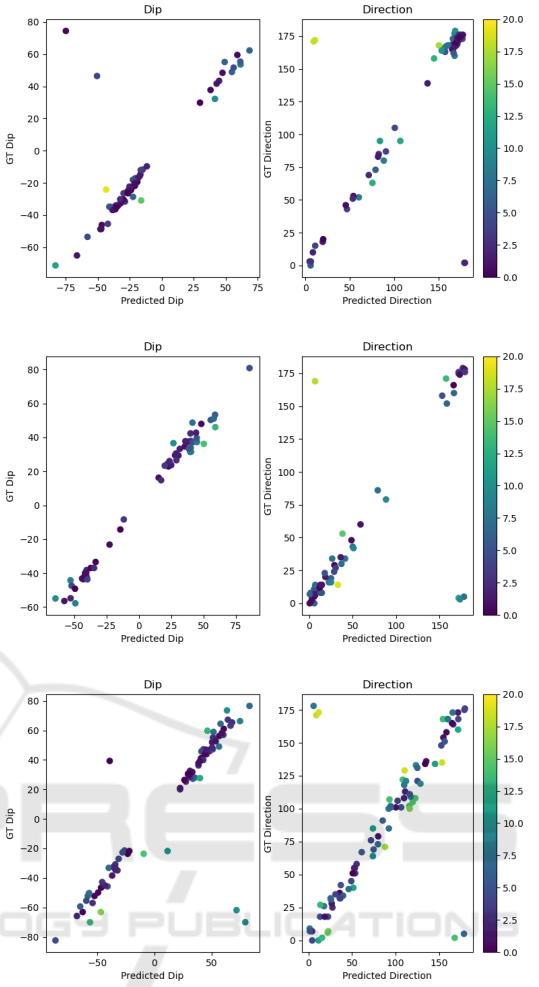

Figure 6: The results of dip and dip direction calculation for the three datasets.

overlaying the detected sinusoid on the ground-truth and counting the common pixels. From a computer vision perspective, it would be enough to classify the detections based on the IoU score, however, given that our end goal for the system is to deliver sinusoid parameters, it is more desirable to compare the detections with the ground-truth in the borehole reference system, namely comparing fractures' depth, dip angle and dip direction by computing the absolute error. In Fig. 5, the PR curve is computed for multiple error thresholds, expressed in absolute units (5-30%). This has been done to overcome the limitation of generating a fixed width for all fractures. However, we can select a decent error margin of 20 units and then calculate the confusion matrix, which can be inspected in Table 1. After carefully analyzing the results depicted in Fig. 6, we came to the same conclusion as (Alzubaidi et al., 2022), that dip direction results are less accurate than dip angle or depth of fracture re-

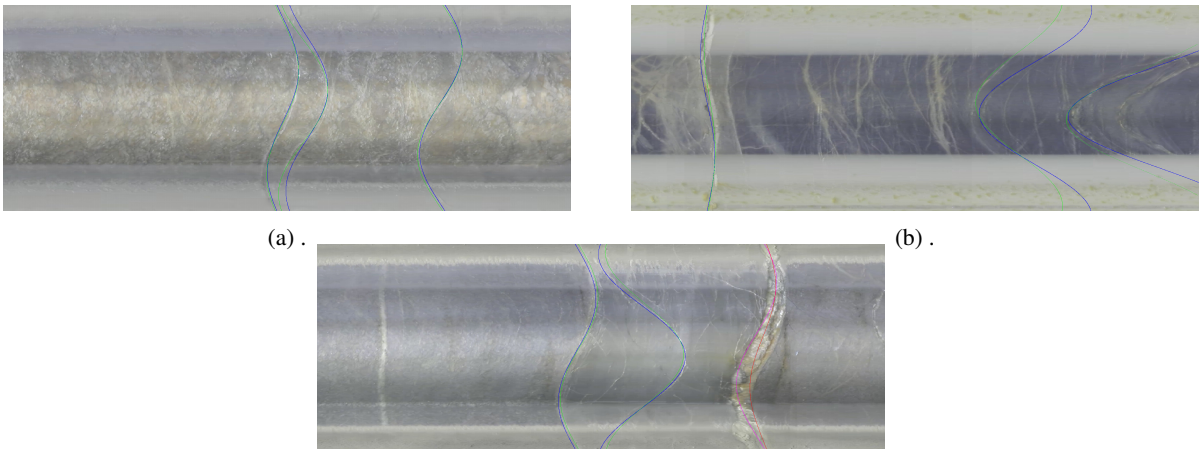

(a) . (b) . (c) .

Figure 7: Sinusoid fitting results. Color scheme: Blue - ground-truth (GT) sinusoid, detected; Green - predicted sinusoid, matched with GT; Red - GT sinusoid, not detected; Pink - predicted sinusoid, not matched with GT.

sults, due to being highly sensitive to the sinusoid fitting and fracture segmentation. Moreover, in the case of the confusion matrix, if one sinusoid prediction is off by more than a certain percentage, then it will produce both a false positive and a false negative. Unfortunately, due to the restricted application domain, there are no open-datasets available to benchmark our method.

Table 1: Confusion matrix for a threshold of 20 units.

Datasets	TP	FP	FN
Dataset#1	66	39	58
Dataset#2	55	17	36
Dataset#3	89	50	89

6 CONCLUSION

This work paves the way towards a fully automatic system for fracture recognition and characterization in borehole images. The proposed workflow trains a semantic segmentation network by using only real images and performs inference to detect possible fracture pixels. The second stage of the workflow fits sinusoid curves on the detected fractures and retains the best matches, ultimately delivering as output the sinusoid parameters in the borehole space parameters, namely fracture depth, dip angle and dip direction. Additionally, we output a confidence score that can be used in further filtering of the detections by a human operator. The developed algorithm is accurate and faster than a human interpretation, hence representing a major help for geologists. As a future work, we aim to benchmark other DNN for semantic segmentation and improve the robustness of the algorithm.

ACKNOWLEDGEMENTS

This work was supported by Real-Time AI-Supported Ore Grade Evaluation for Automated Mining - RAGE project. We also gratefully acknowledge the support of ASTROCK Oy for providing the borehole images and the corresponding ground-truth annotations. The work is part of the Academy of Finland Flagship Programme, Photonics Research and Innovation (PREIN), decision 320168.

REFERENCES

Alzubaidi, F., Makuluni, P., Clark, S. R., Lie, J. E., Mostaghimi, P., and Armstrong, R. T. (2022). Automatic fracture detection and characterization from unwrapped drill-core images using mask r–cnn. *Journal of Petroleum Science and Engineering*, 208:109471.

Anatoli Quintanilla Cruz, R., Cacau, D. C., dos Santos, R. M., Ribeiro Pereira, E. J., Leta, F. R., and Gonzalez Clua, E. (2017). Improving accuracy of automatic fracture detection in borehole images with deep learning and gpus. In *2017 30th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 345–350.

Carey, J. W., Lei, Z., Rougier, E., Mori, H., and Viswanathan, H. (2015). Fracture-permeability behavior of shale. *Journal of Unconventional Oil and Gas Resources*, 11:27–43.

Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. (2015). Semantic image segmentation with deep convolutional nets and fully connected crfs.

In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Chen, L., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y., editors, *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VII*, volume 11211 of *Lecture Notes in Computer Science*, pages 833–851. Springer.

Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1800–1807. IEEE Computer Society.

Dias, L. O., Bom, C. R., Faria, E. L., Valentín, M. B., Correia, M. D., de Albuquerque, M. P., de Albuquerque, M. P., and 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*, 191:107099.

Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html.

Girshick, R. (2015). Fast r-cnn. In *International Conference on Computer Vision (ICCV)*.

Glossop, K., Lisboa, P. J. G., Russell, P. C., Siddans, A., and Jones, G. R. (1999). An implementation of the hough transformation for the identification and labelling of fixed period sinusoidal curves. *Comput. Vis. Image Underst.*, 74:96–100.

Hall, J., Ponzi, M., Gonfalini, M., and Maletti, G. (1996). Automatic Extraction And Characterisation Of Geological Features And Textures Front Borehole Images And Core Photographs. volume All Days of *SPWLA Annual Logging Symposium*. SPWLA-1996-CCC.

He, K., Gkioxari, G., Dollár, P., and Girshick, R. (2017). Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988.

Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q. V., and Adam, H. (2019). Searching for mobilenetv3. In *ICCV*.

Li, J., Jiang, F., Yang, J., Kong, B., Gogate, M., Dashtipour, K., and Hussain, A. (2021). Lane-deeplab: Lane semantic segmentation in automatic driving scenarios for high-definition maps. *Neurocomputing*, 465:15–25.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In Fleet, D., Pajdla, T., Schiele, B., and Tuytelaars, T., editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.

Moran, M. B. H., Cuno, J. S., Riveaux, J. A., Vasconcellos, E. C., Biondi, M., Clua, E. W., Correia, M. D., and Conci, A. (2020). Automatic sinusoidal curves detection in borehole images using the iterated local search algorithm. In *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*, pages 255–260.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*.

Tabernik, D., Šela, S., Skvarč, J., and Skočaj, D. (2019). Segmentation-Based Deep-Learning Approach for Surface-Defect Detection. *Journal of Intelligent Manufacturing*.

Thapa, B., Hughett, P., and Karasaki, K. (1997). Semi-automatic analysis of rock fracture orientations from borehole wall images. *Geophysics*, 62(1):129–137. cited By 32.

van Ginkel, M., Kraaijveld, M. A., van Vliet, L. J., Reding, E. P., Verbeek, P. W., and Lammers, H. J. (2003). Robust curve detection using a radon transform in orientation space. In Bigün, J. and Gustavsson, T., editors, *Image Analysis, 13th Scandinavian Conference, SCIA 2003, Halmstad, Sweden, June 29 - July 2, 2003, Proceedings*, volume 2749 of *Lecture Notes in Computer Science*, pages 125–132. Springer.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.

Wedge, D., Holden, E.-J., Dentith, M., and Spadaccini, N. (2015). Automated structure detection and analysis in televiewer images. *ASEG Extended Abstracts*, 2015(1):1–4.

Zhang, X. and Xiao, X. (2009). Detection of fractures in borehole image. In Zhang, T., Hirsch, B., Cao, Z., and Lu, H., editors, *MIPPR 2009: Automatic Target Recognition and Image Analysis*, volume 7495, pages 1043 – 1048. International Society for Optics and Photonics, SPIE.