

Large volume analysis of core and thin section images in the assessment of Brazil pre-salt reservoir distribution

Edward Jarvis^{1*}, Haoyi Wang¹, Jonathan Dietz¹ and Thomas Van Der Looven¹ discuss how machine learning and artificial intelligence methods can screen a large corpus of unstructured documents, locating and analysing core and thin section images for the purposes of porosity quantification, assignment of core and thin section scale sedimentary facies and the detection of hydrocarbon shows.

Introduction

Understanding the distribution of reservoir intervals in the sub-surface requires the integration of various image and numerical datasets. These take time to locate and use, particularly over large areas and a high number of wells. Core and thin section images are common datasets, typically taken over zones of potential reservoir significance. These images are, however, very qualitative and are therefore under-utilised. Any associated data related to these images, such as core description data, for example, and the samples themselves are also typically generated in a semi-quantitative manner and are therefore slow to generate and prone to descriptor bias (Lokier and Al Junaibi, 2016). These same image-related datasets are also not a direct indicator of reservoir quality, requiring further integration with other datasets before conclusions on reservoir distribution can be drawn.

Technology developments in image analysis techniques have allowed for the prediction of geological features across thousands

of images once a model has been trained. However, while techniques exist and have been applied to various geological image types (Rubo et al., 2019; Falivene et al., 2022; Dietz et al., 2023), it is apparent that fully integrated workflows linking various model results together for further analytical purposes are not as common.

In this article, we discuss how various machine learning and artificial intelligence (AI) tasks have been utilised to efficiently and consistently achieve the following:

- Identification and segmentation of target images from a larger corpus of documents;
- QC of the images to identify those suitable for further analysis;
- Analysis of the images in the quantification of porosity, identification of core-scale sedimentary facies (cm scale), hydrocarbon shows and thin section microfacies (mm scale);
- Integration of the image analysis results across various scales and cross-data validation.

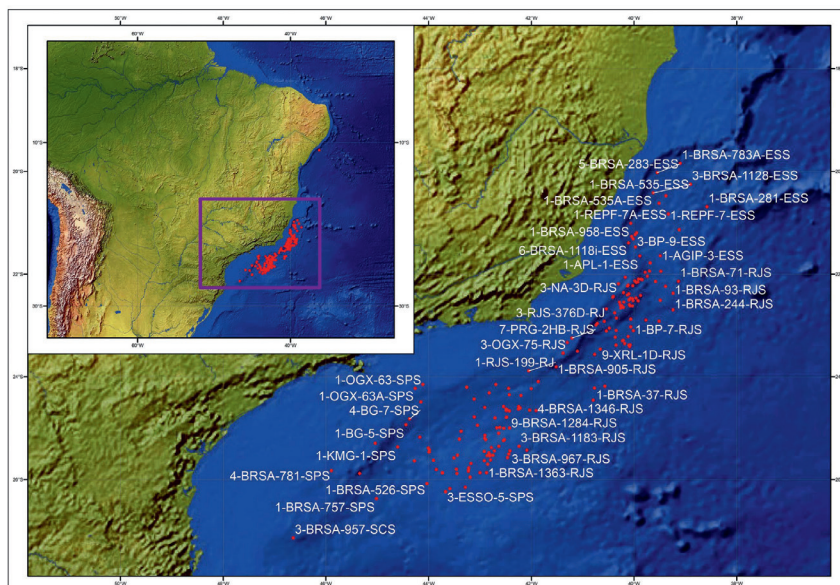


Figure 1 Map of the well locations, Campos and Santos Basins, Brazil.

¹ CGG

* Corresponding author, E-mail: Edward.Jarvis@cgg.com

DOI: 10.3997/1365-2397.fb2024018

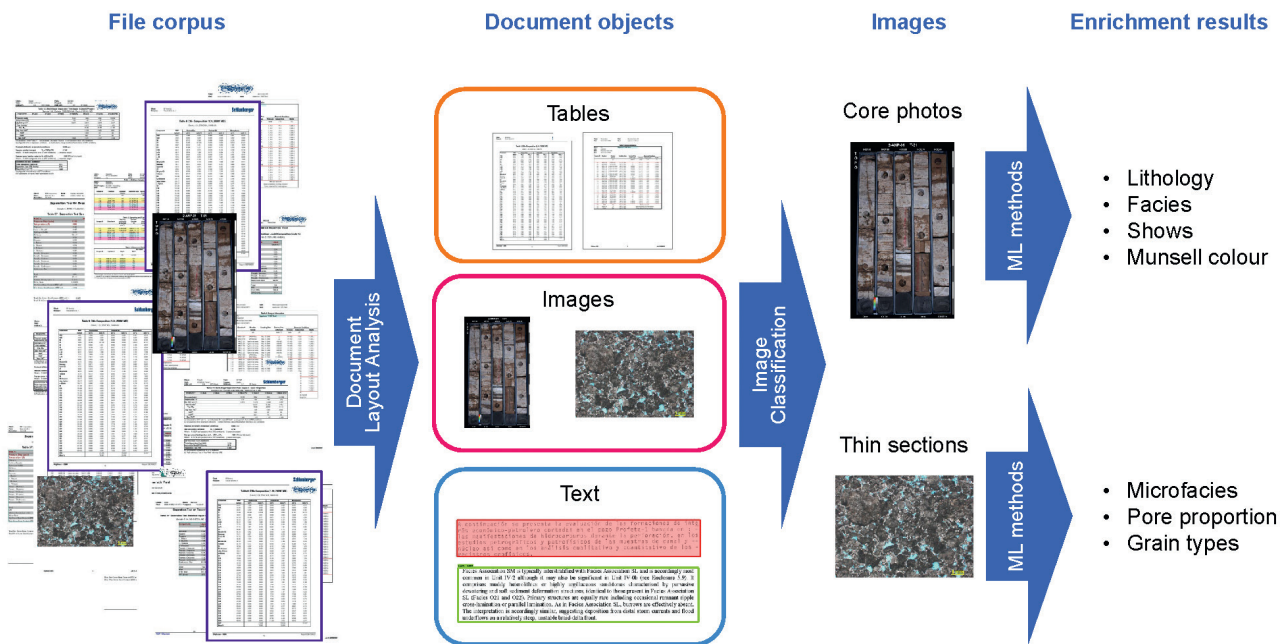


Figure 2 Document classification and image segmentation pipeline.

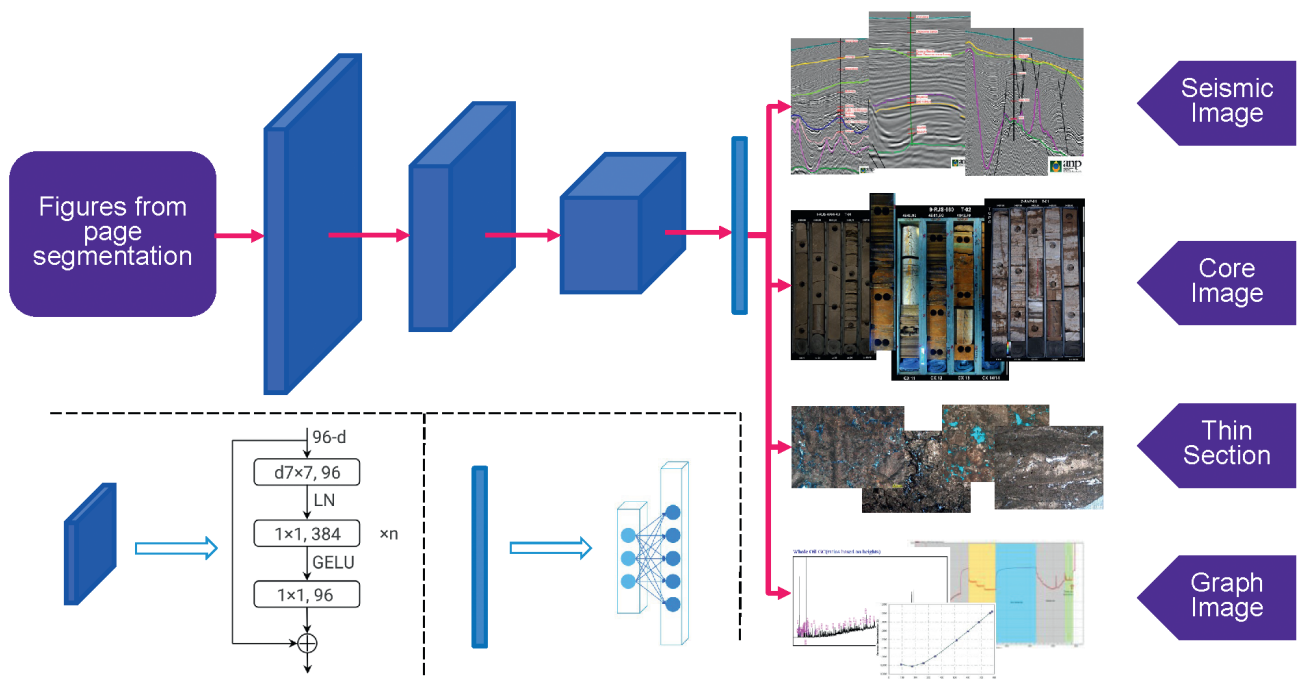


Figure 3 CNN model for the prediction of image classes where each cube represents a stack of ConvNeXt blocks and the rectangle indicates fully connected layers. The actual number of ConvNeXt blocks, i.e., n, used in each stage depends on the architecture of the deployed model.

The process of integrating AI results can be a significant task, particularly as the various models generate large volumes of new data that can be difficult to utilise and time-consuming to relate and compare between the models and across scales. In the increasingly common scenarios that now involve significant automation in data generation and analysis, it is more important than ever to integrate across data types and scales to ensure the results of automated workflows are accurate as well as accessible.

The dataset used in this approach was from the pre-salt stratigraphy of Brazil, with a focus on 228 wells from the Campos and Santos Basins (Figure 1). The data comprised 78,018 docu-

ments in variable formats with a target of locating and utilising the following data for reservoir screening purposes:

- Thin section photomicrographs
- White light and ultra-violet light (UV) core images
- X-Ray Diffraction (XRD) mineralogy
- Core analysis data, such as helium porosity measurements

Data and image identification, extraction and quality screening

An immediate challenge in working with any legacy dataset is the task of locating and extracting the required data from

the source documents. In order to locate the data and images of interest, a first phase of document layout analysis and classification was conducted using a combination of techniques, including computer vision and natural language processing (Lun et al., 2022). Specifically, the authors proposed a directed acyclic graph (DAG) pipeline for automated classification, extraction, and curation of data from various types of documents, significantly enhancing data processing efficiency and accuracy. This approach resulted in the creation of 182,641 document labels, of which 23,494 were images, and enabled the removal of duplicates. Figure 2 illustrates the general workflow and pathways for targeted data types.

A final classification step was then run on the identified images using a convolutional neural network (CNN) (Liu et al., 2022) to identify the required thin section and core images from the wider previously identified image corpus (Figure 3). From the total image dataset, 14,300 thin section images and 1856 core images were identified, with the latter covering 1,247.04 m of vertical section.

The extracted core and thin section images underwent further segmentation to remove background and peripheral content from the images, that if not removed would compromise the image analysis at later stages. File lineage was maintained, linking images to source documents, and any available meta data, such as light source, magnification and depth, was also captured and attributed per image. This final processing step resulted in refined bounding box positions for the 14,300 thin section images and the generation of 6964 individual 1m-long core images that were then stitched into 70 composite images, one per cored interval. Each pixel row was depth-referenced through available meta-data acquired per image.

Prior to any image analysis, a series of image QC steps were conducted to screen the suitability of each image for subsequent analytical steps. This included a further phase of image de-duplication based on hashing approaches, image size and aspect ratio outlier detection, extent of blur/poor resolution based on object boundary sharpness, image brightness and low information/degree of variance (Cleanvision model was utilised for this process: <https://github.com/cleanlab/cleanvision>). From this process, 365 images were identified to be of lower quality. Those images were labelled such that any results, while generated for comparative purposes, could be filtered out of the working datasets at a later stage. Comparison of

results from blurred versus non-blurred images may be useful should Generative Adversarial Network (GAN) (Goodfellow et al., 2014) techniques be applied in the future that could generate new, equivalent synthetic images where the effects of blurring have been removed. A process of image denoising was conducted to sharpen grain/pore boundaries acting to either increase or decrease total porosity values for a given sample. The process was particularly effective in finer-grained samples where boundaries were less well defined even where image resolution was high. In this scenario, the denoising typically reduced the total porosity by a small amount.

Figure 4 illustrates examples of blurred and dark images that were identified in the process and their scores according to the model results. The image quality screening model generally performed well across the image set but some of the finer-grained lithologies or those containing greater abundance of clay were incorrectly categorised as blurred, indicating that further modifications and training were required in certain areas of the model to improve performance.

Thin section images

Thin section images contain significant information on reservoir quality. An example is visual porosity, for which the abundance, size, shape, orientation and distribution of pores can be observed. Further characteristics of the imaged rock can also be determined at this scale, such as the presence and types of cements, grains and clays, all characteristics which infer information on why a rock is or is not porous.

Using image segmentation and object classification on the 14,300 identified thin section photomicrographs, it was possible to rapidly predict proportion and properties of grains, clay, cement and porosity, a task that would have taken many weeks to months to achieve if conducted after manual descriptive techniques. This preliminary result would also allow for further high-resolution models to be trained on specific stratigraphic intervals, reservoir zones and/or geographic areas to produce higher-confidence results.

The first task was to identify pore space in the thin sections and determine the proportion in each image, using image segmentation. The second task focused on identifying grain type and pore space and, as a result, general image classification. Grain segmentation in the determination of grain textural properties was

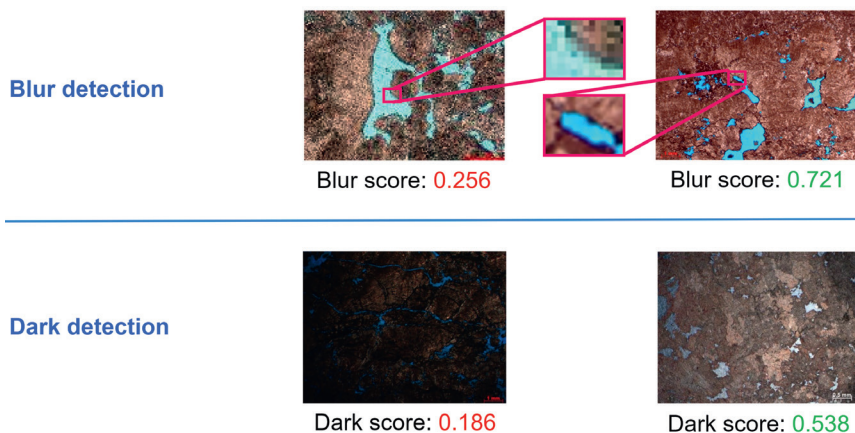


Figure 4 Examples of image quality screening. Scores are on 0 to 1 scale, with a higher score indicative of greater image quality.

not conducted although it is an approach that can be conducted, particularly on elastic rock.

Pore space in thin section images could be discriminated from background regions using colour channels. Since the pore space is a relatively consistent blue under plain polarised light (PPL) and an inconsistent dark blue/black under crossed polarised light (XPL), it was possible firstly to transform RGB images into the HSV colour space in which the hue colour channel controls the visual colour of each pixel. The corresponding hue range was then defined that best represented the pore space, resulting in a mask for each image, as shown with one image example in Figure 5. A filter was set to remove the smallest pores, with a threshold set to remove pores smaller than 5% of the length of the image. This filter could be adjusted by the user and allowed the generation of multiple iterations with different set thresholds. The proportion of pore space for the given image could then be computed by dividing the number of pixels in the pore space region by the number of pixels in the entire image.

A further model was run to identify scale bars. Once identified, optical character recognition (OCR) was undertaken in the extraction of scale-related text, which, in addition to the scale bar, could be used to assign a width in microns for a given image, thereby enabling all pores in the image to be assigned with quantified dimensions.

Post-processing of the segmented pores resulted in statistics for pore count, circumference, length, width, shape, long axis alignment and areal distribution over the slide. Pore dimensions were defined in number of pixels but also microns if a scale bar was present in the image.

Finally, the segmented pore dataset was screened in the identification of outliers and errors using a series of rule logic. Fifty-five rule statements were run over the data resulting in the creation of 1.7 million QC records, flagging either errors or warnings related to individual pores or image and sample level concerns, such as absence of scale bars or no depth datum. The QC process resulted in 835 warnings and errors, identifying pores with very irregular aspect ratios, very large dimensions or regular shapes/straight edges. Such flags could be artificial features,

such as plucked grains, slide edges or sample delamination, and should therefore be removed from further analysis or could be features of interest, such as vugs/ areas of dissolution or natural fractures. One hundred and forty nine images were flagged as high magnification views and/or total porosity values exceeding 45%, likely relating to high magnification images focused on individual large pores. Values derived from high magnification images were flagged, with more representative values taken from lower magnification images, if available.

The applied rules were aggregated to provide an overall rating for a given image and its data, this rating then being used in assessing the image and its data application in various technical use cases, such as an input in petrophysical calibration or general reservoir quality assessment. All QC rules and associated use case terminology and definitions are aligned with the Open Subsurface Data Universe (OSDU) and are therefore universally recognised as an industry standard.

The characterisation of rock samples goes beyond the quantification of porosity, with the assignment of rock facies representing a complex task involving the concatenation of details on porosity, grain and diagenetic phases. For the focus of predicting facies for use in the petrophysical log space, a CNN model was trained in the identification of rock properties, including grain types, clay matrix, cements and, ultimately, microfacies (Dietz et al., 2022). Due to the limited number of samples in the dataset, it was necessary to leverage and fine-tune a pre-trained model to achieve a better classification accuracy. The thin section scans were also divided into smaller image patches to simultaneously increase the number of training samples with a classifier appended with the same number of output neurons to the end of the model. Subject-matter experts initially identified five main rock classes or microfacies for which a training set was generated from 725 images. These microfacies classes were assigned over the training set:

1. Calcareous matrix with grains
2. Mudstone (Carbonate Dunham definition)
3. Porous limestone
4. Cemented
5. Bivalve floatstone

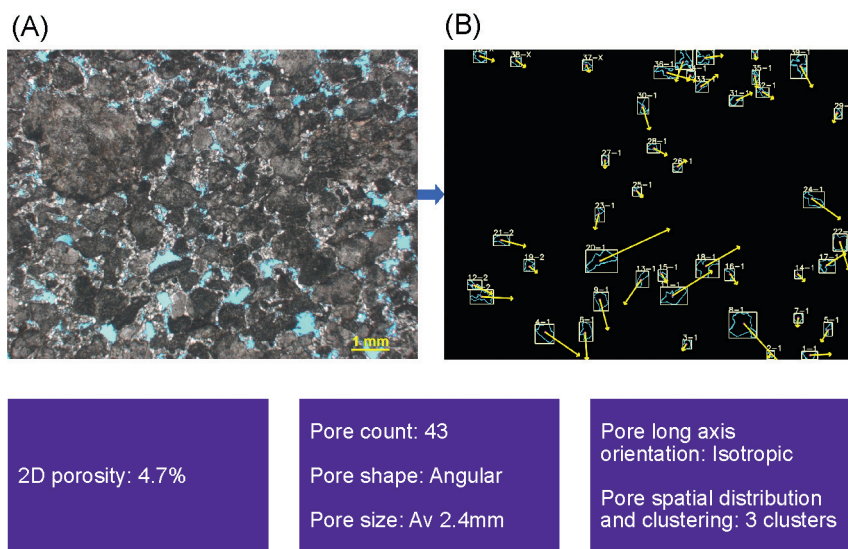


Figure 5 Thin section pore space segmentation. Original image (A) and segmentation mask (B) highlighting individual pores and their long axis orientations as denoted by arrows.

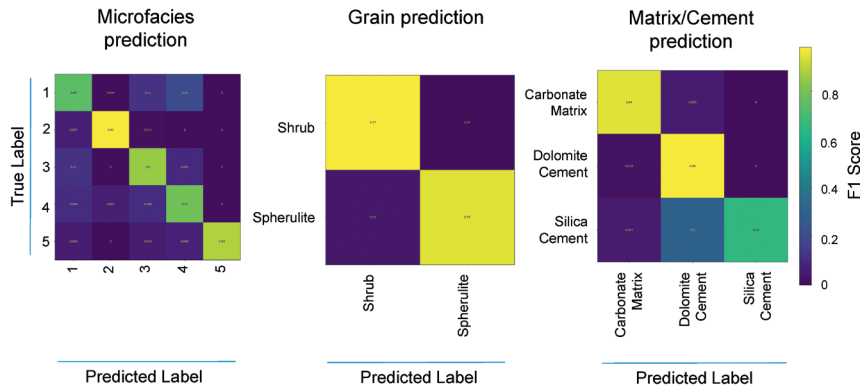


Figure 6 Confusion matrices for thin section microfacies and grain and matrix/cement prediction models. (1) Calcareous matrix with grains, (2) Mudstone, (3) Porous limestone, (4) Cemented and (5) Bivalve floatstone.

A further model was trained in the identification of grain types and surrounding rock matrix. These included spherulites, shrubs, silica and dolomite cement and carbonate clay matrix, which are general terms and definitions associated with the pre-salt stratigraphy and aligning with schemes outlined by Wright and Barnett (2015).

The CNN assigned a predicted microfacies class per image patch in addition to cement, clays and grain classes. Each prediction also included a confidence or F1 score as an indication of the performance of the model. Figure 6 outlines the confusion matrices for thin section microfacies, grain and matrix prediction models. Overall, the model performed well in the identification of most classes. The prediction of cement classes performed well although differentiation of the type of cement was less clear with silica cement frequently being incorrectly predicted as dolomite. This related to a paucity of training labels, particularly for dolomite cement.

Core images

The 70 stitched white light and UV core photographs were analysed with the purpose of identifying features, such as core colour, facies and UV fluorescence. Pixel colour analysis was then conducted on white light images with assignment to the Munsell colour scheme at pixel level. An aggradation step was conducted in selecting the dominant pixel colour per pixel row. Principal component analysis (PCA) indicated common colour clusters relating to different rock properties, such as the presence of clay, oil stain or cement, and therefore provided the means of classifying coarse rock properties or colour-defined facies classes via pixel RGB value.

Rock colouration in the white light images was driven principally by the degree of cementation (resultant white colouration due to low porosity and high silica and or dolomite content), degree of oil stain (pale-orange to dark-brown colouration dependent on porosity distribution and/or associated fluid interaction with formation) and volume of matrix (matrix-rich intervals typically dark-grey to brown due to discolouration/oxidation of carbonates and subtle increase in trace clays and oxides). From these criteria, the subject-matter experts defined four discrete rock classes, namely:

- Non-oil stained, grain-dominated
- Cemented
- Oil-stained or calcareous matrix
- Core plugs and gaps

The training relationship between colour and facies was conducted on a subset of wells which generated a corresponding core image mask indicating the facies assignment per pixel on the photo (Figure 7). A proportion of the wells were kept as a test set while the rest were used in training and validation.

The selected model employed Naïve Bayes to predict the mask directly and, from that, identify the dominant facies (Webb et al., 2010). Naïve Bayes is a probabilistic modelling method that can generate the prediction as well as the probability without a sophisticated training process and can achieve superior performance when there is a strong correlation between the input and the output. The Naïve Bayes is formulated as:

$$P(H|E) = \frac{P(E|H) \times P(H)}{P(E)}$$

where H indicates the facies profile that the method is trying to predict, and E is the evidence on which the prediction is based. In this case, the evidence corresponds to pixel values. $P(H|E)$ is the posterior probability and is a conditional probability, which means the probability of H given E. In this case, it is the probability of the facies to be assigned to a certain pixel value. $P(E|H)$ is called the likelihood, i.e., how different pixel values are associated with each facies type. $P(H)$ is the prior probability, which is the proportion of each facies type in the dataset, and $P(E)$ is called marginal likelihood, which is the proportion of each pixel value. Specifically for pixel-level facies prediction, the above equation can be reformulated as,

$$P(litho|pixel) = \frac{P(pixel|litho) \times P(litho)}{P(pixel)}$$

where $P(pixel)$ and $P(litho)$ are computed based on core photos from a sample well, and $P(pixel|litho)$ is computed based on human labelling. The posterior probability is generated for each pixel. Facies masks were then generated based on the predictions at each pixel, with further aggregation in order to assign a facies to each pixel row and depth increment.

A similar approach was conducted on the UV core images to predict intervals with fluorescence and therefore potential hydrocarbon shows, with the known caveat that dolomite can also fluoresce. By combining the results of the white light and UV core image models it was also possible to further subdivide

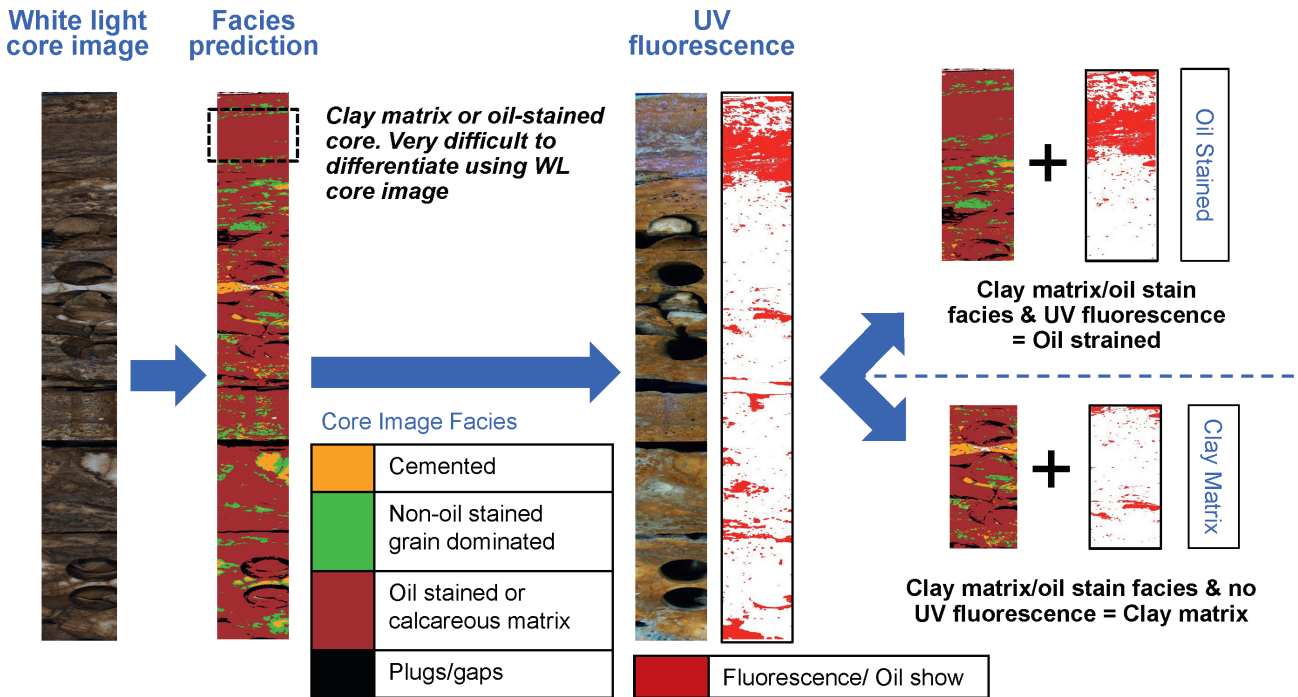


Figure 7 Core image facies and fluorescence prediction.

those classes predicted in the white light model. The presence of a UV fluorescence flag confirmed an oil-stained class prediction over a clay matrix according to the white light model. Figure 7 illustrates the decision path.

Data integration and QC

The thin section and core image analysis models generated predicted facies classes over 1247.04 m of core, 14,300 microfacies and 85,800 grain and cement type occurrences. In addition, over 173,000 individual pores were identified. Six hundred and nineteen thin sections were identified as porous and classified as having ‘good’ reservoir properties based on a porosity threshold of greater than 10%. 907.42 m of potential reservoir facies were predicted over the cored intervals of the study wells, of which 282.27m contained hydrocarbon indicators.

Of the 6070 distinct depths with thin section-derived porosity, 4344 images had a corresponding helium porosity value for cross validation. Interrogation indicated that helium porosity values were on average 7.9% higher than segmented pore values which, in many cases, relates to the difference in measurement precision between the two methods. Helium porosity analysis

is a more precise technique, with means of measuring micropores that are not possible to observe and record in thin section images. Such a differential could be used to quantify the areal percentage of micropores in the samples. Irrespective of the inflated helium porosity values, the overall trends, when plotted on depth, were strongly comparable as illustrated in Figure 8 for one well. An overall conformance of 0.64 cosine similarity was recorded across all wells. After further investigation, samples having a significant divergence beyond a threshold greater than 20% was typically attributed to thin section image capture bias, where results from the images were accurate but were not representative of the entire plug sample. Such scenarios also provide an indication as to the degree of heterogeneity for a given sample.

Further comparison also highlighted a positive relationship between predicted cemented, mudstone and calcareous matrix facies and lower porosities, highlighting a means of qualifying the controls on reservoir quality using the thin section microfacies prediction results. Table 1 illustrates the average porosity and selected XRD-derived mineral weight percentage values for the predicted microfacies classes.

	Predicted Thin Section Microfacies				
	1. Calcareous matrix with grains	2. Mudstone	3. Porous limestone	4. Cemented	5. Bivalve Floatstone
Pore segment (av. area %)	2.54	1.6	8.13	2.59	5.42
Helium porosity (av. volume %)	5.46	7.68	12.75	3.71	5.6
Quartz (av. volume %)	13.92	29.25	10.85	29.58	24.67
Dolomite (av. volume %)	21.04	59	64.76	31.89	37.8

Table 1 Average pore segment, helium porosity and XRD mineralogy values for predicted thin section microfacies.

The results of the cross examination between data sets illustrate trends that would be expected, helping to validate the results of the model. But what is also evident is the demonstrated means of using the predicted microfacies to explain trends in reservoir properties. For all available depth points with thin section imagery, it would be possible to ascertain both the reservoir quality and also the controls on pore volume statistics. While it is possible to derive similar results and observations using manual methods, it is the speed with which the data demonstrated here was generated, from such a wide array of disparate datasets, that highlights the potential significance for subsurface workflows.

The results of QC rule screening and cross-data validation highlighted that a proportion of the segmented pore space data (79.7%) was suitable, according to Open Subsurface Data Universe (OSDU) technical assurance terminology, for general reservoir screening purposes. Suitability was defined where individual pore dimensions and sample scale total porosity values were within acceptable ranges, although sample datum parameters may be absent. 51.6% of the generated data was deemed appropriate for the purpose of formation evaluation calibration where a greater degree of data confidence and precision is required. In this instance, there was a requirement for porosity values between the methods to conform within a 15% differential margin, for clear recorded depth datums to be present, and for pore size attributes and facies predictions to align across both thin section and core images - 20.3% of the image dataset was deemed unsuitable for any further analysis due to one or a combination of factors. Many images failed the image quality pre-processing screening steps, typically being too dark or blurred. Many images were taken at too high a magnification and so were not representative of the samples in general. Porosity extraction and screening indicated some images only contained artificial, oversized or extensively elongated pores through grain plucking or sample delamination, features

typically created during sample preparation and therefore of no value.

When comparing the thin section microfacies prediction results against the core image model, there was conformance between the reservoir classes at both scales in 46.4% of the total instances with conflicts relating typically to localised features only visible at the thin section pore scale, such as silica cements. Figure 9 illustrates such a scenario where two thin sections exist at the same depth with significantly contrasting porosity and microfacies predictions. Comparison with the core image model results and core analysis data indicates both thin section image predictions are accurate, but the lower-porosity image relates to a local cement that also highlights a degree of rock heterogeneity at the centimetre scale. Observations of this kind are quick to ascertain in low sample numbers, but this method of cross-data examination provides a means of quickly screening the heterogeneity of samples and the ability to assess the impact of features observed at the finest scale and their relation to larger-scale reservoir properties.

A final comparison across all the model outputs illustrated a positive correlation between higher UV core fluorescence and typically higher helium porosity values and close correlation with the porous, grain-dominated and oil-stained rock classes as predicted from white light core images and thin sections (Figure 10). Areas of non-conformance seemed to relate to fluorescence occurring over lower-porosity, microporous zones.

Conclusions

This article discusses how the use of AI technologies looked to address the challenge of assessing reservoir quality and its distribution in the subsurface, where the data available for analysis exists in large volumes and unstructured formats. Starting with a corpus of 78,018 unstructured documents from 228 wells, the use of a series of interlinked AI technologies arranged within a pipeline led to identification of the required sub-set of data and

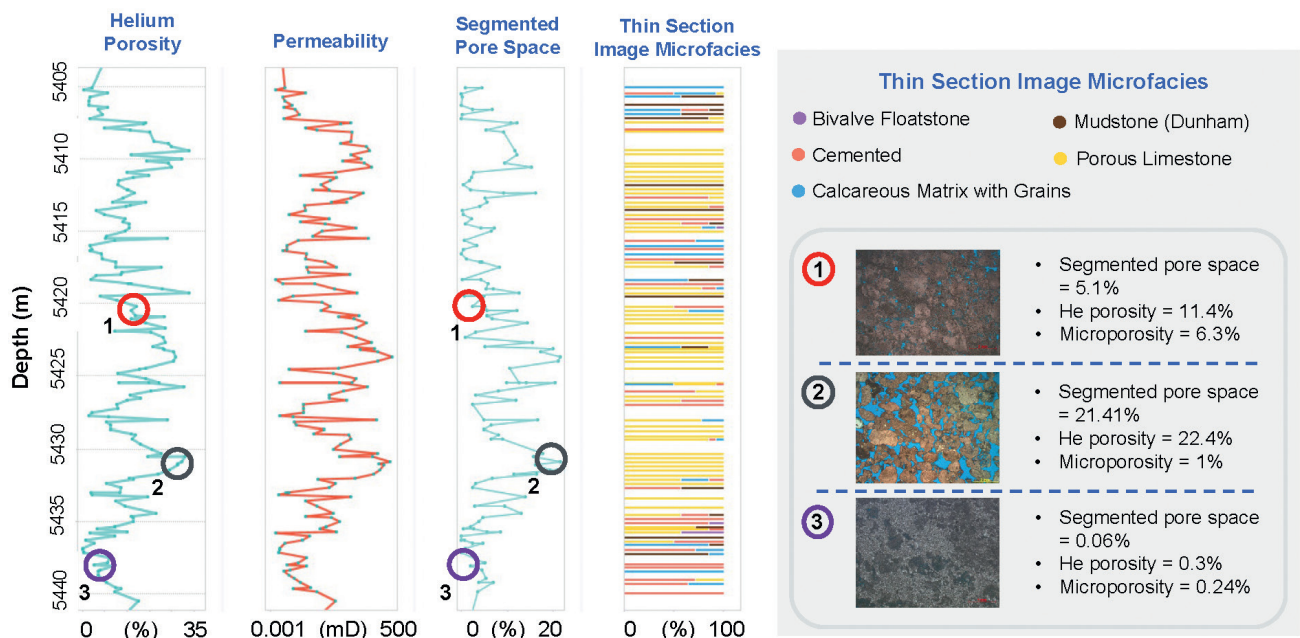


Figure 8 Correlation between thin section, core analysis and thin section microfacies predictions.

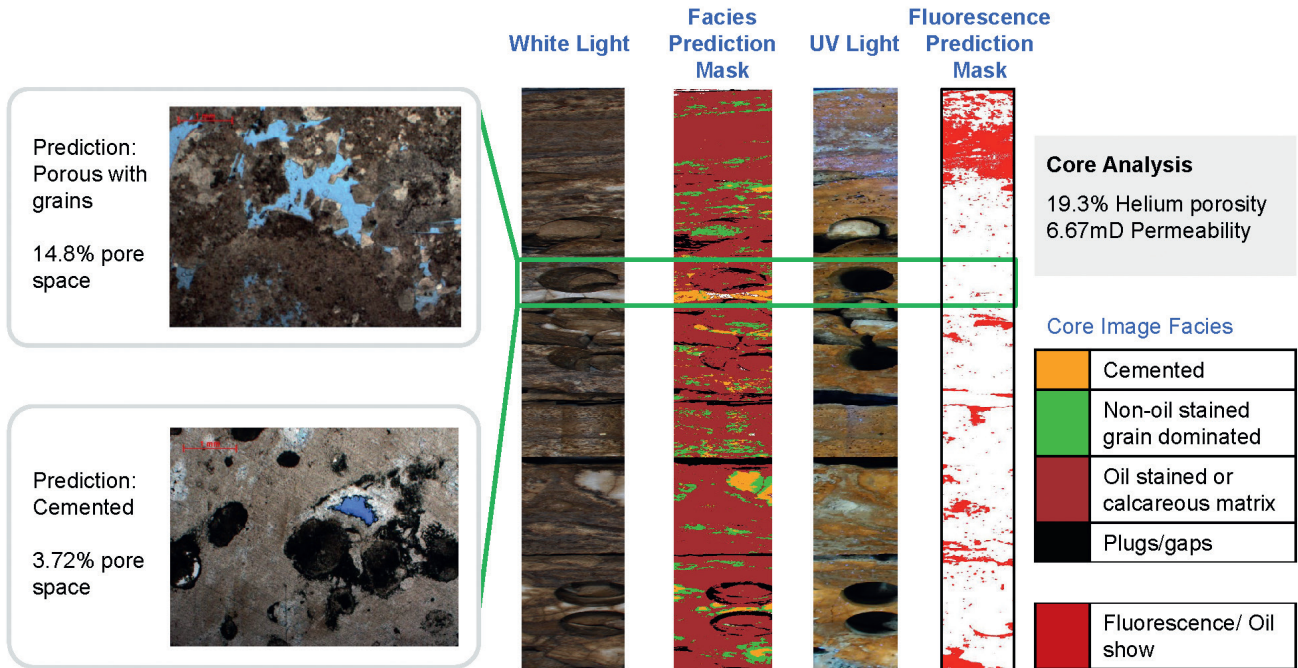


Figure 9 Correlation between thin section and core image model prediction results. A scenario where two thin section images exist at the same depth and have conflicting segmented porosity values reflecting sample heterogeneity. Comparison with the helium porosity measurement indicates which of the images is the most representative.

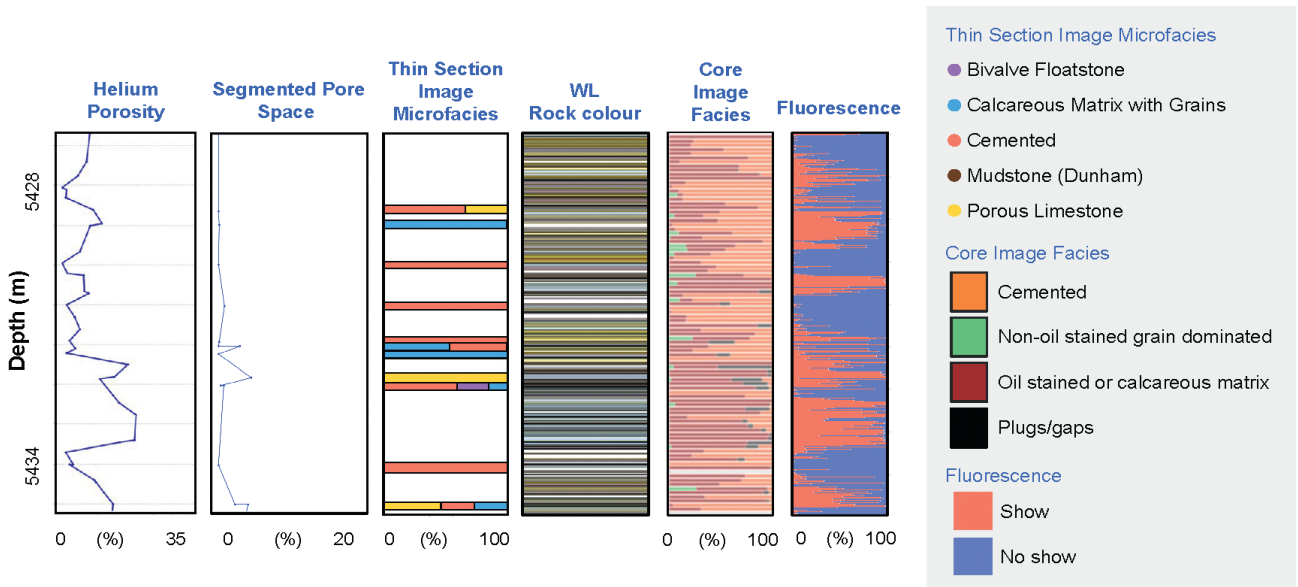


Figure 10 Comparison across models illustrating distribution of UV fluorescence in relation to porosity values and both predicted microfacies and core image facies.

then the generation of new quantitative porosity values and the prediction of microfacies from 14,300 thin section images. A separate process within the same pipeline led to the prediction of geological rock properties and oil show detection from 1856 core photos.

Manual identification, extraction, analysis and integration of these same datasets would be considered a year-long exercise, mobilising various subject-matter experts. In the method outlined in this paper, it was possible to generate porosity values and predict microfacies for the 14,300 thin section images in five days of processing time, equivalent to 30 seconds per image, utilising one CPU and one GPU for inference, once the training model was in place. The processing time for core images was slower, since Naïve Bayes needed to

iterate over all the pixels in the input images, therefore taking dozens of hours to predict the facies for one well. To improve this, deep learning-based image-level models could be developed to reduce the processing time per image from hours to minutes.

The various QC steps and cross-validation flagged errors in the model outputs that could be isolated and also identified geological explanations for variance in reservoir properties. However, it is still apparent that errors and false positives remain to be addressed after later iterations and models. Identifying rock properties in core from colour alone is not sufficient to differentiate between many rock classes, which means that further work is needed to improve the core facies prediction model through the use of a CNN model and labels. This alongside the results

of the current model would help in the identification of textural characteristics in the images. The thin section microfacies model performed well where sufficient labels existed within the training set, and application of the model outside of the Brazil pre-salt stratigraphy indicates a more universal application is possible, with classes such as ‘Mudstone’ and ‘Cemented’ being identified with high accuracy when the model is applied across global carbonate image libraries.

A full screening of the reservoirs discussed in this article would consider additional data types, whether used in traditional workflows or in the methods discussed here. Future work will look to use the model outputs to infill depth sections of the subsurface where core analysis is absent and integrate downhole test data, such as Drill Stem Tests (DSTs) and wireline tool tests, to further validate the results. All results will be used and tested as additional input calibration points to image log and petrophysical interpretation models.

In summary, the merits of the approach outlined in this article aim to demonstrate that such methods are a means of both efficiently extracting data from documents and generating large amounts of new geological data in timeframes that would not be possible using manual methods. As the results indicate, the models will never predict results with 100% accuracy. However, the models enable large volumes of data to be screened quickly, with QC flags and certainty scores in place to guide subject-matter experts to those wells of interest or needing further attention and targeted manual assessment. The approach outlined is intended to be a method utilised alongside manual methods to help tackle large datasets, enabling initial higher-level screening across the entirety of regional datasets prior to the implementation of manual techniques in specific areas once a greater regional understanding of data coverage and trends is available.

References

- Dietz, J., Wang, H., Jarvis, E. and Hou, S. [2022]. Unlocking the Potential in Your Core, *Thin Section, and Image Log Data Through Image Processing*. Third EAGE Conference on Pre Salt Reservoirs. European Association of Geoscientists & Engineers, 1-7.
- Dietz, J., Wang, H., Hou, S., Jarvis, E. and Sekti, R. [2023]. *Borehole Image Logs: New Approaches to Automated Surface, Breakout and Facies Interpretation*. Third EAGE Digitalization Conference and Exhibition. European Association of Geoscientists & Engineers, 1-5.
- Falivene, O., Auchter, N.C., Pires de Lima, R., Kleipool, L., Solum, J.G., Zarian, P., Clark, R.W. and Espejo, I. [2022]. Lithofacies identification in cores using deep learning segmentation and the role of geoscientists: Turbidite deposits (Gulf of Mexico and North Sea). *AAPG Bulletin*, **106**(7), 1357-1372.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. [2014]. Generative adversarial nets. *Advances in neural information processing systems*, **27**, 2672-2680.
- Liu, Z., Mao, H., Wu, C., Feichtenhofer, C., Darrell, T. and Xie, S. [2022]. A convnet for the 2020s. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Lokier, S.W. and Al Junaibi, M. [2016]. The petrographic description of carbonate facies: are we all speaking the same language? *Sedimentology*, **63**, 1843-1885.
- Lun, C. H., Hewitt, T. and Hou, S. [2022]. A machine learning pipeline for document extraction. *First Break* **40**(2), 73-78.
- Rubo, R.A., Carneiro, C.C., Michelon, M.F. and Gioria, R.S. [2019] Digital petrography: Mineralogy and porosity identification using machine learning algorithms in petrographic thin section images. *Journal of Petroleum Science and Engineering*, **183**.
- Webb, G. I., Keogh, E. and Miikkulainen, R. [2010]. Naïve Bayes. *Encyclopedia of machine learning* **15**(1), 713-714.
- Wright, P. and Barnett, A.J. [2015]. An abiotic model for the development of textures in some South Atlantic early Cretaceous lacustrine carbonates. *Geological Society, London, Special Publications*, **418**.