Whole core scanning and artificial intelligence for the automatic recognition of sedimentological features

Christophe Germay[1], *Tanguy* Lhomme[1,*], *Luc* Perneder[1], *Bastien* Pinchard[1], *Thomas* Mazur[1]

[1]EPSLOG, Liege, Belgium

Abstract. CoreDNA is a sustainable approach to core analysis, developed to optimize the value of cores by increasing the quantity of high-quality data extracted while keeping the analyst footprint to a minimum. CoreDNA combines a suite of cost-efficient, non-destructive tests to produce a multi-disciplinary dataset, as early as possible in the core analysis workflow, without causing any irreparable damage to the cores. Using a fully mobile technology, core analysts generate continuous high-resolution logs of rock properties, immediately upon barrel-opening, prior to irreversible core alteration by slabbing and plugging. The same technology can be successfully mobilised on whole cores, all formats of legacy cores, but also plug samples and to some extent, on cuttings. Tracks ranging from textural and colour indices measured on the rock surface, grain size distribution statistics, elemental concentrations, elastic wave velocities and rock strength are generated. This data is integrated under a unique format and used as a roadmap to guide core analysts during more complex and expensive discipline-specific tasks such as SCAL or RMT. CoreDNA data is suitable to feed high-density multidisciplinary databases with input data formatted to train high-end predictive models for geological and petrophysical descriptions using the latest developments in Machine Learning. The CoreDNA workflow also benefits from a unique digital core visualization platform that enables detailed sedimentological core descriptions based on high and ultra-high-resolution photographs of cores. It provides hugely important information and data on the fabric and texture of reservoir rocks, and input into the interpretive elements of core description such as depositional process and environment assignments. Since compositional, textural, structural, and diagenetic features of the rock are concisely annotated in depth-referenced ultra-high-resolution photographs, this objective information is adequate to train instance segmentation models for the automatic recognition of sedimentological features over extensive core depth intervals. With an iterative, multistage approach to Artificial Intelligence (AI) model building and training in mind, we developed graphic tools to visualize and quickly perform quality checks on the automated recognition of sedimentological features such as grains, laminations, fractures, fossils, etc. This helps validate the accuracy of the AI algorithms and identify areas where further improvements can be made. We describe these functionalities in detail and use real-world examples ranging from reservoir characterization and optimization in hydrocarbon-bearing formations to quantitative resource mapping for base mineral deposit management to demonstrate how these help materialize our vision of sustainable core analysis.

1 Introduction

This paper is the third instalment in a series of publications on the combination of high-resolution core data and artificial intelligence at the edge of the ongoing elevation of core analysis standards. In the first paper of this series [1], we have unravelled concepts involving our state of the art technology for the acquisition of high-resolution core data and how such data could be fed to unsupervised machine learning schemes for the identification of a number of facies utilising a cluster analysis which establishes 'groupings' of data points with similar physical, visual and elemental characteristics. In a second publication [2] we have shown how recent developments in Deep Machine Learning (DML) schemes could be combined in an iterative, multi-stage and cost-conscious approach to core analysis where complex rock properties, which normally requires expensive testing protocols, are predicted by trained Artificial Intelligence (AI) models, fed with inexpensive high-quality data obtained from transdisciplinary, high-resolution, non-destructive measurements on whole cores.

In this present paper we focus on the tools specifically developed to make this iterative approach usable in the context of real-world applications of core analysis for the technical evaluation of subsurface resources, while keeping in mind the importance of maintaining large and well-groomed databases wired to streams of high-fidelity data, thoroughly checked for quality before being used for the training of performant predictive AIs. The overarching architecture of this work is depicted in Figure 1, in which

* Corresponding author: tanguy.lhomme@epslog.com

the logic of interaction between the different components and the underlying data links are symbolized. Examples of applications of this workflow to solve technical challenges from the extractive industry are presented in the later sections of this paper.

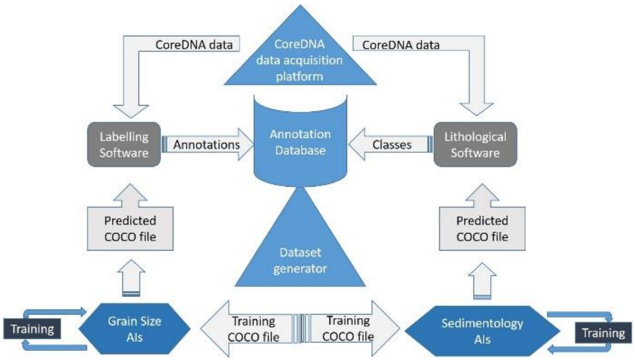

Figure 1: CoreDNA workflow using Artificial Intelligence models for the prediction of grain size distribution and lithofacies.

Software and tools

This section describes the elements symbolized in Figure 1 and their role in a systematic iterative approach towards the building of performant predictive AIs.

1.1 CoreDNA data acquisition

CoreDNA is an integrated core analysis solution combining transdisciplinary, high-resolution, non-destructive measurements on whole cores, for an early yet objective description of cores and the rapid estimation of formation properties. Details on the CoreDNA data acquisition platform are given in [1].

Results of these fast tests are analysed real-time and turned into high-resolution, continuous profiles of properties (petrophysical, geomechanical and geochemistry).

1.2 Lithological software

LithoLog is a software module developed with input from expert sedimentologists, with the goal of digitalising sedimentological information using standardized qualifiers to minimise subjectivity and inconsistency in the core description process, while also generating robust, fully quantitative sedimentological data at a scale normally reserved for detailed petrographic analysis/optical microscopy. Since its previous version described in [2], the software has been transformed from a standalone version towards a seamless integration into the architecture presented in Figure 1. The interfacing between the LithoLog and the annotation database system was implemented by standardizing the annotations file format and organizing the project folder tree.

1.3 Labelling software

The labelling software module is a CoreDNA data visualisation platform equipped with manual and AI-assisted annotation creation and edition tools. These tools have been designed for the task of object segmentation in the high (35µm/pixel) and ultra-high (1.8µm/pixel) definition pictures that are standard components of CoreDNA datasets. Custom lists of user-defined labels are created on a per-project basis, while polygonal annotations are created and edited with a manual drawing tool (Figure 2) or with the following AI-assisted object segmentation tools.

1. The AI is tasked with the segmentation of one single object in the whole image. This mode is useful for pre-processing (third-party) images of multiple cores before they're fed to subsequent segmentation AI searching for detailed sedimentological features.
2. The AI is tasked with the segmentation of one single object in the area delimited with the rectangle selection tool (Figure 3).
3. The AI is tasked with the segmentation of as many objects as possible in the area delimited with the rectangle selection tool
4. The AI is tasked with the segmentation of as many objects as possible in the entire picture.

The labelling software module is designed with maximum modularity in mind, so that fit-for-purpose AIs can easily be selected in adequation with the segmentation task at hand. The modular architecture of the labelling software stems from the use of a specific Application Programming Interface (API), which sets the rules used to communicate via TCP between its different components:

- A National Instrument Labview graphical user interface (GUI)
- A Python-based backend server running in the background, waiting for orders from the GUI to run one of the three tools described above.

Labels are used for the creation of compact annotation datasets stored in the central annotation database. Such datasets are suitable for the training of classification or segmentation AIs. Although the labelling software is designed to exploit the specific attributes of CoreDNA data (length, the project folder tree structure and its well/core organisation, and various user defined configuration files), it can also be used with third party image data streams.

Figure 2: Labelling software main panel.

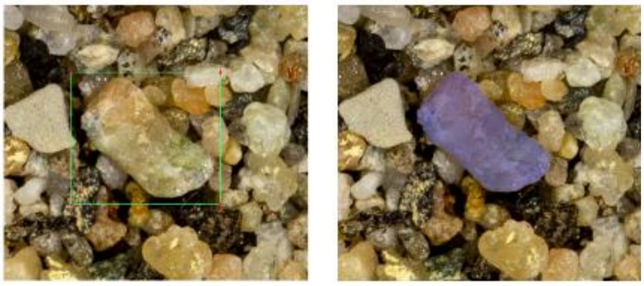

Figure 3: Example of the AI-assisted annotation creation tool. The Region of Interest (ROI) defined with the rectangle selection tool is shown in the left-hand-side while the resulting AI-assisted object segmentation is visible in the right-hand-side image.

1.4 Annotation database

The annotation database system is the main storage component for labelled CoreDNA data. It is designed to receive data from the labelling and lithological software modules and to treat queries from the dataset generator. The two aforementioned pieces of software create and add tags to each annotation in order to indicate their potential usefulness for a variety of available AI models. Being centimetric classification labels or pixel-wise segmentation polygons, both annotation types use similar protocols to be stored into the annotation database. This annotation database is a conceptual implementation of a central database, meaning that no conventional architecture was used.

The conceptual database consists in effect in a series of dedicated file naming and writing conventions. The annotation files are stored in project subfolders in accordance with the pre-existing CoreDNA data storage files and folders structure. The annotation file queries issued by end-users are handled by the dataset generator described in the next paragraph. These annotation files are also accessible with any conventional file manager. Furthermore, annotations can easily be retrieved, displayed and exported together with other relevant CoreDNA data subsets.

1.5 Dataset generator

The dataset generator module is the link between the annotation database and the AI models that are used for making predictions with technical applications such as object detection, segmentation, and classification in high-resolution core photographs. The purpose of the dataset generator is to receive and automatically process queries for the selection, assembly and transformation of bundles of annotations stored in the annotation database into Common Object in Context (COCO) files suitable for the training of user-specified AIs for specific tasks. The COCO dataset is one of the most popular large-scale labelled image datasets available for public use. It is widely understood by state-of-the-art neural networks and defines an adequate data format suitable for standard AI tasks. This standard dataset format is vastly supported by Python DML developments and was thus the best choice for defining the dataset generator output format. COCO stores data in a JSON file formatted by information, licenses, categories, images, and annotations. This JSON formatting was initially introduced for instance segmentation dataset but can easily be used for object detection, (sub-)image classification, etc. The dataset generator software thus sends queries to the annotation database system to retrieve and consolidate datasets for specific AI applications by generating a COCO formatted file. Moreover, seeing that annotated data can be created by multiple users throughout multiple projects and with various objectives in mind, this module should be capable of remapping the labels from any selected project (possibly created by multiple users) to a new custom list of unified annotation classes that are wanted in the final dataset (Figure 4). This last feature enables multi-party collaborations and data gathering efforts spanning multiple projects without having to predefine and enforce the use of a global list of annotations labels.

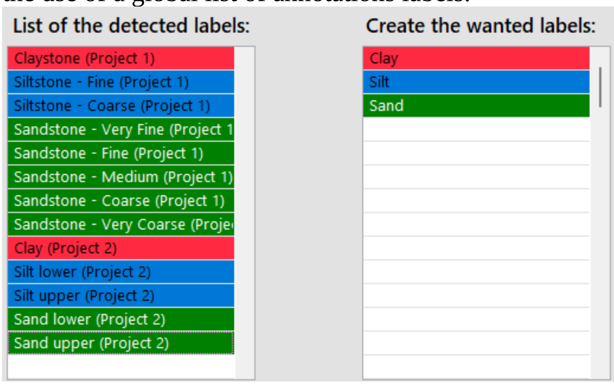

Figure 4: Handling and unifying class labels in a multiple-user collaborative environment.

2 AI models

2.1 Classification

Classification in the field of AI applied to computer vision refers to the task of categorizing or labelling objects or images into predefined classes or categories. The goal of classification is to train a machine learning model that can automatically assign the correct class label to unseen images based on the patterns and features it has learned from a training dataset [3].

2.2 Segmentation

In computer science and image processing, segmentation refers to the process of dividing an image into meaningful regions or segments. This technique is used to identify and extract specific objects or regions of interest from an image. There are several standard segmentation techniques in the field of AI and computer vision (Figure 5, [4]), which specificities are addressed below

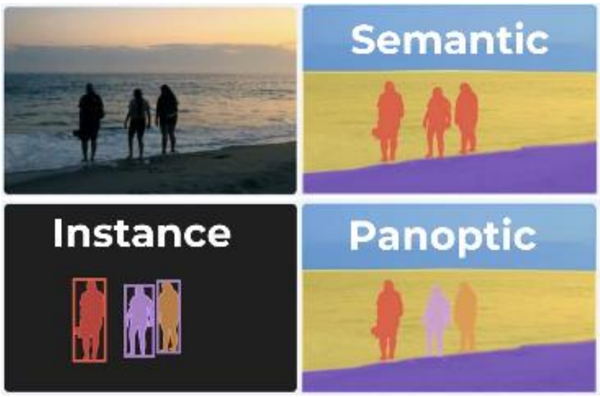

Figure 5: Comparison of the three presented image segmentation method. The original image is the top left.

Semantic segmentation

Semantic segmentation focuses on pixel-level classification, where each pixel is assigned to a label representing the object or region it belongs to. It involves deep learning models, such as convolutional neural networks (CNNs), to extract and analyse image features hierarchically, capturing local patterns and context and using this information in order to segment images and assign a class to each of their pixels. In the context of core analysis, semantic segmentation can help with the automatic classification of core intervals in predefined lithofacies (shales, sandstone, carbonate, etc.).

Instance segmentation

The goal of instance segmentation is to identify and classify individual objects within an image, typically by localizing them with bounding boxes and defining their contours by labelling each instance separately. In the context of this paper, an instance segmentation AI assisted by DML algorithms was developed and trained for the task of isolating and sizing individual grains in ultra-high-resolution photographs of core samples and cuttings from clastic formations. Sizing of sub-micrometre scale objects in these UHR photographs is made possible by the knowledge of scale at a pixel level. It could also be applied to the automatic detection of carbonate micro-textures, as studied in [5]. The automatic detection of laminations, fractures, fossils, in core images add value to sedimentological studies by saving a significant amount of time while bringing consistency and accuracy in the analysis process, especially when dealing with large datasets involving multiple wells.

Panoptic segmentation

Panoptic segmentation extends these concepts by unifying instance and semantic segmentation. It aims to partition an image into a set of non-overlapping regions, where each region is associated with a specific object instance and labelled with a class category. The result is a pixel-level segmentation map that distinguishes between different object instances and provides semantic labels for each region. For instance, panoptic segmentation can be applied to the same use-cases than semantic or instance segmentations techniques while having a higher level of detail. For example, it can be used to derive mineral maps from core photographs by leveraging its ability to segment and classify objects within an image, or to detect and locate multiple cores from pictures of large core trays containing several core pieces.

3 Applications

3.1 AI models for the automatic classification of lithofacies

In a previous instalment in this series of papers [2], we've shown how CoreDNA data could be fed to a supervised machine learning algorithm using CNNs for the automated identification of lithofacies, the design of fit-for-purpose plug selections and the programming of subsequent steps in core analysis programs. Recent progresses made since this first demonstrator of study include the automatization of this task through the integration of the annotation database system and the dataset generator software.

3.2 Instance segmentation of individual grains from an oil and gas bearing formation

In this section we describe an application of the instance segmentation logic described in section 2.2 to the task of deriving a continuous profile of percentiles from grain size distributions mapped with centimetre resolution in ultra-high-resolution photographs of core samples from a clastic reservoir.

The performances of several approaches involving AI models were evaluated against a "ground truth" obtained with a different kind of method. The ground truth for the grain size distribution consists in the median grain size derived from the analysis of topography maps produced with a laser beam scanning the MiniSlab surface continuously along whole cores, with a vertical accuracy of 1 micron and a horizontal resolution of 20*20 microns This method was checked against sieve tests and gave positive results for grain sizes above 15µm (see [1] for details).

Prior to using instance segmentation techniques, we tried to use an end-to-end approach that performed regression to predict values of the 50th percentile (P50) of the grain size distribution from series of one centimetre wide slices of the original CoreDNA images. The model was trained using the laser topography analysis results as ground truth. The results of this approach are shown for a small dataset in Figure 6. In this case, we can see that smaller grain sizes predicted by the AI compare well with the ground truth obtained from the laser topography analysis. However, the difference between both methods increases for larger grain sizes.

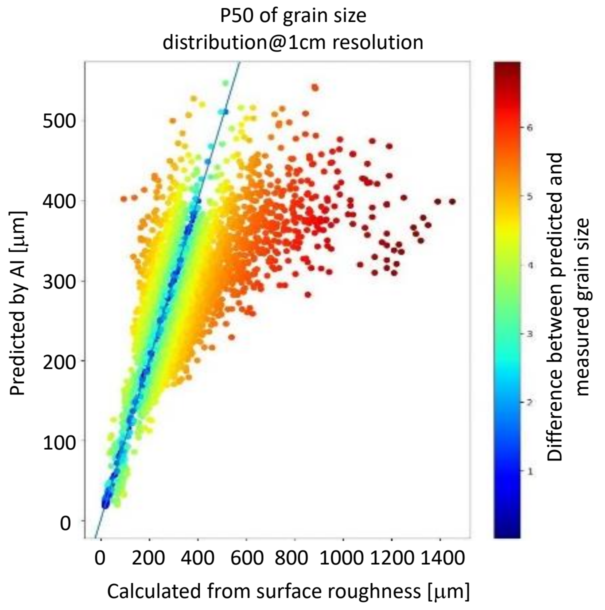

Figure 6 : Scatter plot of performances of an end-to-end model predictions compared to the ground truth (calculated from surface roughness measurements).

A major inconvenience with this kind of method comes from the absence of intermediate results. This entails a low explainability of the model, which in turns makes iterative model refinements and performance improvements difficult, if not limited.

In a second attempt, we tried Mask R-CNN [6] as a proof of concept for the feasibility of the instance segmentation approach. Mask R-CNN was judged adequate for this task, as the flagship architecture to tackle instance segmentation problems while it is also quick to get running. With only a small dataset of 30 images and about 9000 instances, it proved the feasibility of the intermediate step of segmenting the grains as an intermediate result. This conclusion was established on the comparison of AI predictions with two independent methods for the evaluation of grain size distributions: a laser topography and a Wentworth bin classification [7] based on a core description performed by an experienced sedimentologist. In Figure 7 we show the profile of the median of the grain size distribution derived from the analysis of the laser topography of the core surface, with a centimetre resolution. The colour scale for the data points shown in this plot represents the difference in µm between the P50 estimates from the laser measurement and the middle value of the bins selected by the sedimentologist for the corresponding 1cm interval. In Figure 8 we show the profile of median grain sizes calculated for the distribution of grains segmented by the AI model on 1cm strips of the original UHR photographs of the core sample. The colour scale for the data points shown in this plot represents the number of grains segmented in the corresponding strip. In both figures the light blue rectangles represent the Wentworth classification given by the expert.

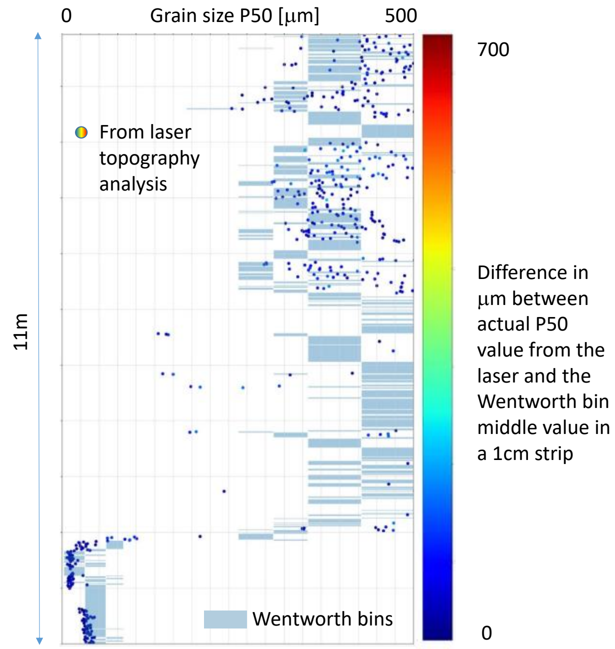

Figure 7: Profile of the median grain size estimated from the laser topography map. The light blue rectangles represent the classification in Wentworth bins.

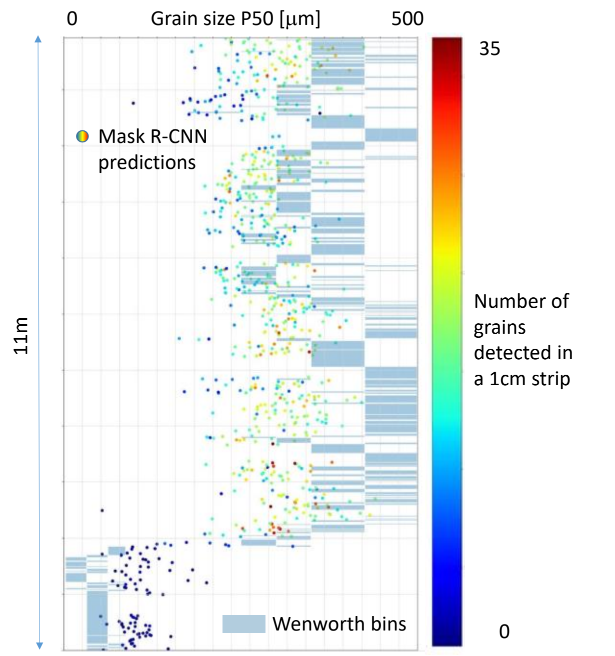

Figure 8: Profile of the median grain size measured on the segmentation of the grain by Mask R-CNN. The light blue rectangles represent the classification in Wentworth bins

Although these results can be deemed encouraging in view of the very small size of the training data set available at the time, we decided not to pursue the work with Mask R-CNN because of limitations of this approach in terms of resolution of the input images and of the number of segmented objects per image. The set requirement for the model to detect grains in a window spanning 1cm in length meant that the model must be able to handle 5300x4000 pixels. At the time of this trial, feeding such a large a picture to a Mask R-CNN model was not feasible so we had to resort to the following methods in order to be able to process the data.

Downscaling: The first method consists in downscaling the image so that smaller versions can be fed to the model. The problem with this approach is that it removes the finer details in the picture and tends to bias the model toward bigger grains (Figure 9).

Tiling: The second method consists in splitting the image into smaller pieces, or tiles (Figure 10). In this case, a grain spanning two adjacent tiles could be missed or wrongly measured. This edge effect happens more frequently as the grains are bigger compared to the size of the tiles, which tends to bias the model toward smaller grains. Nevertheless, it can be mitigated by allowing adjacent tiles to overlap.

Figure 9 : grain segmentation with Mask R-CNN on a downscaled photograph of cuttings, courtesy of Stratum Reservoir.

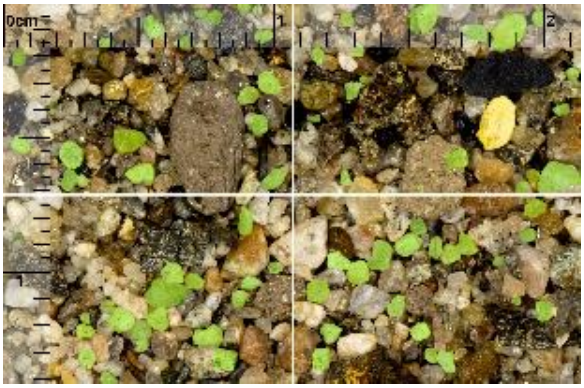

Figure 10: grain segmentation with Mask R-CNN of the full scale picture covering the same area with four adjacent tiles of a picture of cuttings, courtesy of Stratum Reservoir.

These two methods can be combined to create unbiased predictions for large, high-resolution images, at the cost of a significant increase of the prediction time. Moreover, their usage duplicate the number of times a region of the picture is processed, which is time-consuming and requires a post-processing algorithm to removes grains that have been segmented multiple times.

Furthermore, the architecture of Mask R-CNN imposes the declaration of a maximum number of instances to segment from a given picture a-priori. Hardware limitations fixed this amount to 100 instances per picture in our case. Even though we managed to artificially increase the limit with the tiling explained above, we chose to change the interface of our model to get rid of this limitation.

In a third and latest approach, we used the labelling software and its embedded segmentation capabilities detailed in section 1.3, in order to identify and contour grains in a semi-automatic manner. We then tested the performance of our approach by computing the median grain size on the segmentation results and comparing it to the analysis of topography as shown in Figure 11.

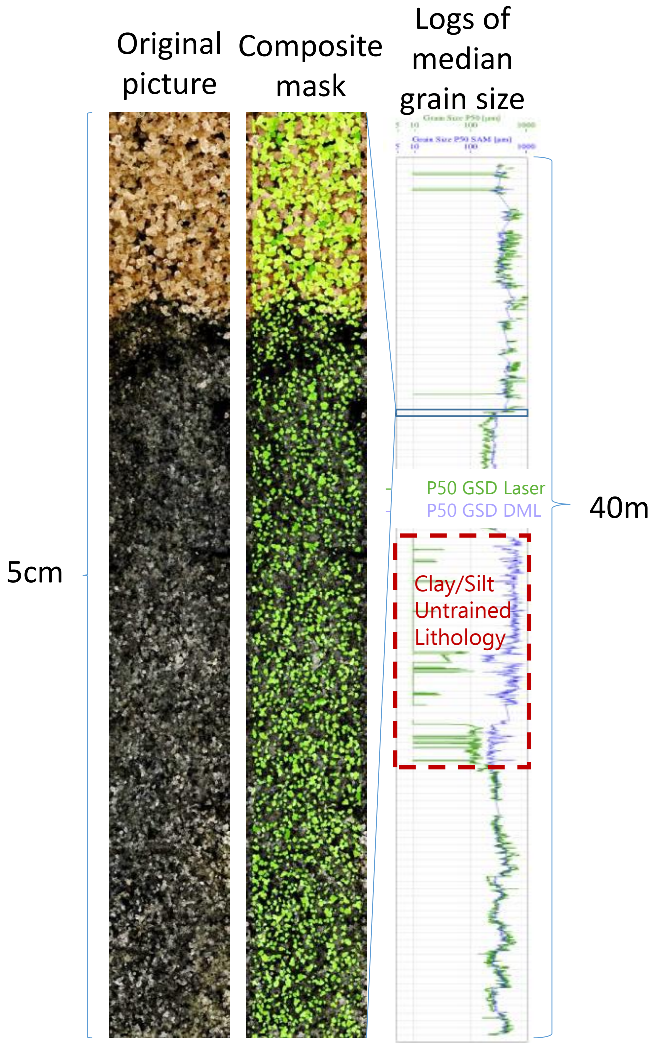

Figure 11 : Comparison of the median grain size profiles obtained with a segmentation AI and with the analysis of laser topography maps

The performance are overall rather good, although they strongly depend on lithology. The Clay/Silt interval appearing in the middle section of Figure 11 is characterized by a very poor match between the P50 predicted by the segmentation AI and the ground truth. This is a clear sign that a method built around the segmentation of individual grains from high resolution photographs would hit a limit when the boundary of individual grains can no longer be seen. Such a limit has not been quantified yet, and progress remains possible in terms of maximal image resolution that the segmentation logic could handle. If this resolution happens to be too low to give useful results in shale rocks, other models involving an end-to-end approach (such as the one presented earlier in this section) could be used to tackle rocks beyond this limit.

3.3 Mining

In this section we describe a case study for the identification of Garnet mineral grains from ultra-high resolution (UHR) pictures of drill cores. A semantic segmentation AI was trained on Quantitative Evaluation of Minerals by Scanning Electron Microscopy (QEMSCAM) analysis 2D maps from rock samples for which UHR photographs were also available. SEM-derived automated mineralogy was acquired using the QEMSCAN® platform collected by Rocktype Ltd. The scanning electron microscope (SEM) has backscatter electron (BSE) and energy dispersive X-ray (EDS) detectors to provide automated petrographic quantification of geological samples in the form of spatially resolved compositional and textural data. In this study, all data were collected using a QEMSCAN® WellSite instrument (Quanta 650 F FEG Scanning Electron Microscope) using the FieldScan mode at 15 kV beam energy and 5 µm or 10 µm step intervals. The mineral data were processed in FEI's iDiscover software package using an in-house mineral library developed by Rocktype Ltd. The UHR photos were analysed with a DML model using a U-Net shape with an EfficientNet [8] backbone.

A transfer learning strategy was used and a customized training procedure was developed by freezing specific parts of the model weights arrays during training. The dataset was acquired on ore-rich cores, studied, and then complemented with various cherry-picked images containing important visual features. A Comparison between the QEMSCAN mineral map and AI predictions for Garnet are shown in Figure 12.

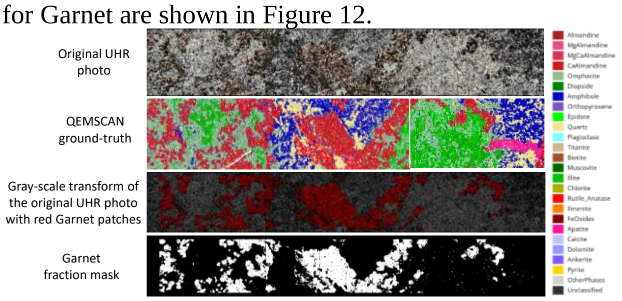

Figure 12: Comparison of the QEMSCAM training data (second picture from the top) and of results obtained by the garnet detector AI on a UHR photograph (first picture from the top of an ore-rich sample. The mineral of interest (Garnet) is highlighted in red in the third picture from the top. The corresponding mask is shown in the bottom picture.

The trained Garnet detector was used to predict the mineral abundance on multiple wells with various visual characteristics.

4 Conclusions

We have developed a complete framework dedicated to the exploitation of the latest advances in Artificial Intelligence for the purpose of modernising core analysis by taking full advantage of the volume of high-quality, high-resolution data produced with the CoreDNA multi-sensor core testing platform. A collection of proprietary software modules have been developed and interlinked in a compact and seamless architecture to enable the fast and objective-driven analysis of large datasets from industrial projects for the extractive industry.

This development represents a new step towards a more sustainable approach to core analysis that maximizes the value of core samples while minimizing the workload of analysts.

Real-world examples ranging from hydrocarbon reservoir characterization and optimization to quantitative resource mapping for base mineral deposit management demonstrate the practical applications and tangible benefits of CoreDNA in achieving sustainable core analysis. By combining advanced technology, comprehensive data acquisition, and AI-driven analysis, CoreDNA realizes the vision of sustainable core analysis, enabling more efficient and informed decision-making in various industries.

References

[1] C. Germay, T. Lhomme and P. Bisset, "Combining high-resolution core data with unsupervised machine learning schemes for the identification of rock types and the prediction of reservoir quality.," in *The 34th International Symposium of the Society of Core Analysts*, 2021.

[2] C. Germay, T. Lhomme and L. Perneder, "Combining high-resolution core data and machine learning schemes to develop sustainable core analysis practices," in *The 35th International Symposium of the Society of Core Analysts*, Austin, 2022.

[3] F.-F. Li, J. Johnson and S. Yeung, *Convolutional Neural Networks for Visual Recognition*, Stanford University, 2017.

[4] N. Rogge, S. Singh and A. Dirik, "Universal Image Segmentation with Mask2Former and OneFormer," 2023. [Online]. Available: https://huggingface.co/blog/mask2former.

[5] C. Birnie and V. Chandra, "Using deep learning for automatic detection and segmentation of carbonate microtextures," in *SEG/AAPG International Meeting for Applied Geoscience & Energy,*, Houston, 2022.

[6] K. He, G. Gkioxari, P. Dollar and R. Girshick, "Mask R-CNN," in *IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, 2017.

[7] C. K. Wentworth, "A Scale of Grade and Class Terms for Clastic Sediments," *The Journal of Geology*, 1922.

[8] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Conventional Neural Networks," in *International Conference on Machine Learning*, Long Beach, 2019.

[9] C. Shorten and T. Khoshgoftaar, "A survey on Image Data Augmentation for Deep Learning," *Journal of Big Data*, vol. 6, no. 60, 2019.