
Detection of Geobodies in 3D Seismic using Unsupervised Machine Learning

Supervisors:

Author:

Quentin Corlay

Prof. Vasily DEMYANOV

Dr. Dave McCARTHY

Dr. Dan ARNOLD

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

Institute of GeoEnergy Engineering

School of Energy, Geoscience, Infrastructure and Society

Heriot-Watt University

May 2023



The copyright in this thesis is owned by the author. Any quotation from the thesis or use of any of the information contained in it must acknowledge this thesis as the source of the quotation or information.

Abstract

In this work, we present a novel, automated method for detecting geobodies in 3D seismic reflection data, helping to reduce interpreter bias and speed up seismic interpretation. A seismic geobody refers to a geometrical, structural, or stratigraphic feature, such as a channel, turbidite fan, or igneous intrusion. Geobodies are subtle seismic features, hard to pick, and their detection is challenging to automate due to their complex 3D geomorphology and diversity of shapes. Nevertheless, the detection and delineation of these structures are essential for improving the understanding of the subsurface as well as building a variety of conceptual models.

In our approach, we can rapidly interpret large 3D seismic volumes using point cloud-based segmentation to identify geobodies of interest, including complex stratigraphic features like lobes and channels. By converting the 3D seismic cube into a 3D seismic point cloud (sparse cube), we reduce the volume of data to analyse, which in turn speeds up the detection process. First, we build the 3D point clouds by filtering the seismic reflection volume using different seismic attributes, and then each point in the cloud is segmented into different clusters. The clustering is performed using the unsupervised Density-Based Spatial Clustering of Applications with Noise (DBSCAN) which allows the segmentation of all structures present into delineated objects. The clustered objects can then be characterised by features based on their 3D shape and spatial amplitude distribution. Finally, our method allows the selection of a specific geobody and can retrieve geobodies based on their similarity to exploration targets of interest.

The method has been applied successfully to two modern 3D seismic datasets (Falkland Basins) and two types of geobodies: fans and sill intrusions. We demonstrate that our method can scan through a large 3D seismic volume and automatically retrieve likely fan and sill geobodies in a very efficient manner. This approach can be used to scan through large volumes of 3D seismic, looking for a wide variety of geobodies.

Acknowledgements

I would like to take this opportunity to express my deepest gratitude and appreciation to all the remarkable individuals who have played a pivotal role in my PhD thesis.

To my PhD supervisors. Vasily, thank you for your indispensable guidance, always putting me on the path to solving the problems we were facing, and I wish to particularly thank you for your constant presence at every step of this journey, at every conference or difficulty I encountered, you were there to reassure me and give visibility to my work. Dave, thank you for helping me understand and frame everything about this thesis, making sure we were working on real problems, and looking for practical solutions. Also, for your vigilance making sure every event was celebrated properly. Dan, thank you for the advice you have given me over the years, always getting to the heart of the matter and provoking fruitful thoughts and discussions.

To my coworkers and fellow researchers at Heriot-Watt University, you have been truly inspiring. Special mention to all the present and former members of our GeoDatascience group Bastian, Athos, Chao, Tom, Rhona, Eduardo, Diana, Bertrand, Iris and Farah. Your friendship made this PhD a lively experience and lead to moments that I will cherish. Working with you guys has been a daily pleasure.

I would like to extend my appreciation to the CDT in Oil & Gas, put together by John Underhill, for supporting this work. Thank you to Lorna and Anna for organising memorable field trips and training courses. Being part of a cohort has been a real benefit to me helping me throughout the struggles of a PhD. Denis, Ramy, Luigi, Michael and Robert thank you for sharing this journey with me and making sure we visited every nightclub all over UK and Spain.

I would also like to acknowledge my jury Uisdean and Guillaume and all the reviewers that helped me refine this research work, the discussion we had has been invaluable.

To administrative support and academic staff, I am grateful to all the individuals who have facilitated my research.

To my friends, I recall the memories in Edinburgh, Bergerac, Contis, Greece, Corsica... Your friendship has brightened my daily life and our apero has allowed me to think about something other than my research during these years. Also, I always deliver. To be continued.

To my parents, I owe a debt of gratitude that words cannot fully express. The completion of this thesis stands as a testament to your unwavering belief in me. To my siblings, yes I finally did it! Your encouragement has been my motivation, I hope now you will read it to the very last word.

To Omblin, your support has been the bedrock of my journey. You stood by me during my moments of self-doubt, celebrated my successes, and provided the strength and resilience needed to overcome any challenge and conquer any obstacle. Your craziness has been my escape, your tenderness my refuge and your strength my inspiration. I am extremely proud to be by your side. I am looking forward to the future we will embrace together.

This PhD study was funded by NERC National Productivity Investment Fund (NPIF) grant and Heriot-Watt University via its James Watt Scholarship scheme. Their support is gratefully acknowledged. I also thank the Falkland Islands Government for providing the 3D seismic reflection dataset that forms the basis of this study.

Inclusion of Published Works Form

Please note you are only required to complete this form if your thesis contains published works. If this is the case, please include this form within your thesis before submission.

Declaration

This thesis contains one or more multi-author published works. I hereby declare that the contributions of each author to these publications is as follows:

Citation details	Corlay, Q., Demyanov, V., McCarthy, D., & Arnold, D. (2020, December). Turbidite Fan Interpretation in 3D Seismic Data by Point Cloud Segmentation Using Machine Learning. In <i>EAGE 2020 Annual Conference & Exhibition Online</i> (Vol. 2020, No. 1, pp. 1-5). European Association of Geoscientists & Engineers.
Author 1	Contribution to method development in discussions, code implementation, wrote manuscripts
Author 2	Supervision, contribution to method development in discussions
Author 3	Supervision, contribution to method development in discussions
Author 4	Supervision, contribution to method development in discussions
Signature:	
Date:	24/05/2023

Contents

Abstract	i
Acknowledgements	ii
Table of Contents	v
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Definitions	1
1.2 Research motivations	2
1.3 Statement of the problem	4
1.3.1 Geobodies and their detection in seismic data	4
1.3.2 Challenges	7
1.3.2.1 Challenges in geobody detection	8
1.3.2.2 Challenges in automation	10
1.4 Research methodology	11
1.5 Thesis outline	14
2 Literature review – Automation in seismic interpretation	18
2.1 Interpretation of horizons from 3D seismic data	19
2.1.1 Picking horizons	19
2.1.2 Horizon auto-trackers	19
2.1.3 Global seismic interpretation methods	20
2.2 Machine learning and deep learning in seismic interpretation	22
2.2.1 Definitions	22
2.2.2 Machine learning applied to seismic interpretation	23
2.2.3 Deep learning applied to seismic interpretation	24
2.2.4 Training dataset for seismic interpretation	25
2.2.4.1 Challenges with building a representative seismic dataset	25
2.2.4.2 Open-source data initiatives	26
2.2.4.3 Data augmentation and transfer learning	27

2.2.4.4	Synthetic seismic data	27
2.3	State of the art on geobody extraction	29
2.4	Uncertainty in seismic interpretation	31
2.5	Discussion	33
3	Full Seismic Segmentation with Point Cloud Representation and Density Clustering	35
3.1	Introduction	35
3.1.1	Objectives and requirements	35
3.1.2	Overview of existing methods	37
3.1.3	Our novel method	39
3.2	Point Cloud Seismic	40
3.2.1	Why working with a seismic data point cloud?	40
3.2.2	Creation of a seismic data point cloud	41
3.2.3	Application of the point cloud extraction to the <i>validation dataset</i>	44
3.2.4	Advantages and Disadvantages of a seismic data point cloud	46
3.2.5	Visualisation and plotting of 3D Seismic as a Point Cloud	49
3.3	Segmentation of a seismic data point cloud	51
3.3.1	Approaches to point cloud segmentation in other industries	51
3.3.2	Density-based Segmentation	52
3.3.3	Application of the segmentation to the <i>validation dataset</i> point cloud	61
3.3.4	Evaluation of the Segmentation	64
3.3.4.1	Methods	64
3.3.4.2	Validation on the <i>validation dataset</i>	72
3.3.5	Optimisation of the Segmentation	75
3.3.5.1	Methods	75
3.3.5.2	Validation of the optimisation	78
3.4	Conclusion and discussion	80
4	Detection of 3D geobodies in a segmented seismic	83
4.1	Introduction	83
4.1.1	Seismic attributes for geobody detection	84
4.1.2	Deep learning approaches applied to geobody detection	85
4.1.2.1	State of the art results	85
4.1.2.2	Challenges of deep learning applied to turbidite fan detection	86
4.2	Overview of our approach	88
4.2.1	Object recognition in segmented point cloud	88
4.2.2	Proposed workflow for the detection of geobodies among point cloud objects	90
4.3	Feature engineering and data analysis	91
4.3.1	Geometric shape features	92
4.3.2	Amplitude-based features	96
4.3.3	Feature selection and preprocessing	98

4.3.4	Data analysis	98
4.3.4.1	Univariate and bivariate analysis	100
4.3.4.2	Multivariate analysis - dimension reduction	103
4.3.4.3	Summary of observations from the data analysis	104
4.4	Geobody detection by closest object retrieval	105
4.4.1	Methodology	105
4.4.1.1	Initial attempts at the geobody detection	105
4.4.1.2	Closest object retrieval paradigm	107
4.4.1.3	Binary hashing for fast object retrieval with LSH	108
4.4.2	Results	110
4.4.3	Assessment of the methodology	110
4.5	Conclusion and discussion	113
5	Geobody detection in 3D seismic - Application to real case studies	115
5.1	Introduction	115
5.2	Application to turbidite fan detection in the complete North Falkland Basin dataset	116
5.2.1	Data and methodology	116
5.2.2	Results	118
5.2.2.1	Seismic segmentation	118
5.2.2.2	Fan detection	122
5.2.3	Assessment of the application to the NFB dataset	126
5.3	Application to sill detection in the Falkland Plateau dataset	127
5.3.1	Data and methodology	127
5.3.2	Results	129
5.3.2.1	Seismic segmentation	129
5.3.2.2	Sill detection	131
5.3.3	Assessment of the application to the FPB dataset	133
5.4	Conclusion and discussion	134
6	Conclusion, discussion and future work	135
6.1	Conclusion	135
6.2	Discussion	137
6.3	Future work suggestions	139
A	Code available	142
	Bibliography	143

List of Tables

3.1	Global segmentation scores of the 7 realisations – validation dataset . . .	72
4.1	Eigen-based features	94
4.2	List of the five most common Haralick features	97
5.1	Processing times for the segmentation of the NFB dataset	117
5.2	Processing times for the segmentation of the FISA dataset, FPB	129

List of Figures

1.1	Facies model for a submarine turbidite fan	5
1.2	Types of magmatic bodies	6
1.3	Saucer-shaped sill geometry on a seismic section	7
1.4	The challenge of the detection of turbidite fans from section views	9
1.5	Schematic diagram showing the fan geometries that result from various sediment types	10
1.6	Presentation of the validation dataset and the test datasets.	13
1.7	Geobody detection workflow: from 3D seismic to geobodies	17
2.1	Sketch of a simple auto-picking algorithm	19
2.2	Systematic horizon extraction from global interpretation methods	20
2.3	Workflow computing RGT model by establishing links among elementary horizon patches	21
2.4	Unsupervised and supervised machine learning applied to seismic interpretation	23
2.5	Different ways to label seismic data for machine learning tasks	26
2.6	Workflow to incorporate real information in synthetic data at the training stage and synthetic information in real data at the inference stage	28
2.7	Body delineation from seismic attribute cube	29
2.8	Comparison of the output obtained from four different methods to geobody extraction	31
2.9	Examples of interpretations from a single 2D seismic image classified by tectonic settings	32
3.1	Seismic segmentation for object detection	36
3.2	Workflow of the seismic segmentation method	40
3.3	Extraction of local maxima on seismic wiggle traces.	42
3.4	Creation of a Point Cloud Seismic of a seismic slice	42
3.5	Differences among the coherence implementations	43
3.6	Impact of the amplitude cut-off values on a point cloud seismic slice	45
3.7	Illustration from seismic to point cloud	46
3.8	Visualisation of the point cloud seismic, focused on the turbidite fans	47
3.9	Point Cloud Seismic result of a small 3D seismic cube	48
3.10	Octree structure which subdivides spatially the point cloud into different Level-of-Detail	50
3.11	Level of Details visualisation example	50

3.12	Illustration of the planar segments produced for building interior classification	51
3.13	How DBSCAN works	53
3.14	Vertical exaggeration of a seismic data point cloud	55
3.15	Vertical exaggeration impact on the segmentation	55
3.16	Choice of DBSCAN hyper-parameters	56
3.17	HDBSCAN dendrogram construction	57
3.18	How HDBSCAN works	59
3.19	DBSCAN vs HDBSCAN results on a small point cloud	59
3.20	Illustration of the notion of hierarchy in the geologic structures	60
3.21	Results of the segmentation with DBSCAN	62
3.22	Results of the segmentation focused on the turbidite fans	63
3.23	Comparison between the same turbidite fan extracted manually and by the point cloud segmentation	63
3.24	Main segmentation problems	64
3.25	Horizontally Superposed Layers index approximation	68
3.26	Cluster amplitude entropy vs cluster Zeboudj’s internal disparity	70
3.27	High and low Entropy cluster vs high and low Zeboudj’s internal disparity clusters	71
3.28	7 realisations of the segmentation of the <i>validation dataset</i>	72
3.29	Variation of the global segmentation metrics - <i>validation dataset</i>	73
3.30	<i>HSL</i> and <i>CI</i> per cluster to identify cluster with overlap or heterogeneous amplitude distribution	74
3.31	Flowchart of the optimisation of the segmentation by merge of neighbouring segments	76
3.32	Flowchart of the optimisation of the segmentation by division of under-segmented segments	77
3.33	Sea Lion / Casper fan under-segmentation problem	79
3.34	Optimisation of the segmentation of a cluster	79
3.35	Horizons tracked across fault blocks using Dynamic Time Wrapping for signal correlation	81
4.1	Spectral decomposition and RMS to better image channel definition	85
4.2	Flowchart of training a CNN for channel detection on synthetic seismic before applying it to field seismic	86
4.3	Turbidite fan deposit modelled with stratigraphic forward modelling	88
4.4	Reference point cloud objects and labels from the <i>Mian dataset</i>	89
4.5	Illustration of the diversity of shapes and geometries between turbidite fans	90
4.6	Object detection workflow from a segmented seismic	91
4.7	Workflow to extract the contours of an object in a 2D projection	93
4.8	Comparison of $r_{contours}$ between a horizon and a turbidite fan	93
4.9	Diagram of point cloud eigenvectors and eigenvalues	94
4.10	Illustration of the depth distribution of a horizon and a fan	95
4.11	Visualisation of the main objects present in the <i>validation dataset</i>	99

4.12	Box-plot of the selected features for fan retrieval	100
4.13	Cross-plot matrix of the selected features for fan retrieval	101
4.14	Cross-plots contour ratio vs amplitude mean and contour ratio vs sphericity	102
4.15	3D visualisation of steep slope objects	102
4.16	t-SNE representation of the seven extracted features in a 2D projection .	104
4.17	Illustration of locality hashing of the attribute space and retrieval of the closest object	109
4.18	Results of the 30 most similar objects to a fan on the <i>validation dataset</i> .	111
4.19	Result comparison of the retrieved objects inputting three different fans as an example	112
4.20	Results of the 30-closest objects retrieved from the elongated fan	112
4.21	Encoder-decoder network structure	114
5.1	Map view and cross-section of the NFB seismic dataset	117
5.2	Filtering of small clusters on the NFB dataset	118
5.3	Results of the seismic segmentation of the entire NFB dataset	120
5.4	Fault splitting an object and resulting in a segmentation error	121
5.5	Comparison between the Sea Lion fan delineated from the segmentation of the <i>validation dataset</i> and of the entire NFB seismic survey	121
5.6	Largest cluster in the segmentation displaying a delta system and fan bodies close to the delta	123
5.7	Delineation of the three known fans displayed above the deltaic deposits	124
5.8	Results of the fan detection from the NFB seismic survey	125
5.9	Location of the FISA seismic survey	128
5.10	Sills and lava flow distribution and associated forced-folds	128
5.11	Cropped sequence from the FISA seismic cube, FPB	129
5.12	Results of the point cloud segmentation of the FPB dataset	130
5.13	Vertical seismic section crossing the central sill reflection, FPB	131
5.14	3D visualisation of the shallow sill provided as an example	131
5.15	Sill detection results - 100 closest objects to a sill example	132
5.16	Focus on the main object retrieved, delineation of the central sill and dykes	132
5.17	Focus on the objects retrieved at shallow depths, delineation of small sills with saucer-shaped bodies.	133
6.1	Summary of the methodology presented in this thesis	136
6.2	Immersive Visualisation Environment for geoscience data	139

Chapter 1

Introduction

The primary motivation for this project is to **automatically detect and extract geobodies** from 3D seismic reflection data where geobodies are small, complex stratigraphic features such as lobes, dykes or sills that are particularly important in assessing geoenery reservoirs. We aim to find a fast and generalisable solution so it can be applied to a range of seismic interpretation tasks in any geological setting, combining machine learning with geological knowledge to ensure meaningful features are extracted from 3D seismic reflection data.

1.1 Definitions

Seismic reflection data is one of the principal tools for imaging the subsurface, particularly for applications such as hydrocarbon exploration and resource management. Seismic reflection surveying essentially involves sending sound waves into the Earth's subsurface and collecting the echoes to estimate the subsurface's properties using the principles of seismology and wave propagation. Seismic waves can travel long distances through the Earth, and, as they travel through a medium, they may be refracted or reflected when they encounter a boundary or property change in the medium. Thus, seismic data provides information on the boundaries and changes of fluid content in

the environment traversed over kilometres underground [Eitzenberger, 2012, Fjær et al., 2008, Musgrave, 1967].

The seismic data exploration process consists of three main distinct stages: (i) data acquisition – recording of a seismic survey with a source and some receivers (geophones or hydrophones); (ii) seismic imaging – data processing transforming the seismic records into an accurate image of the subsurface at a given resolution; and (iii) seismic interpretation – inferring meaningful geology from the seismic images based on geological knowledge and understanding.

Geobody extraction is an essential element of (iii) – seismic interpretation and is critical for estimating seismic stratigraphic architectures, defining facies and improving reservoir understanding. Seismic interpretation is often subdivided into three interrelated categories: *structural*, *stratigraphic*, and *lithologic interpretation* [Jarvis and Saussus, 2009], but geobody extraction has implications for all three categories and provides quantitative geomorphology measurements and discrete objects that can be integrated into a geological model.

A geobody (or geo-object) refers to a 3D geological element in the subsurface defined by distinct geometric, structural or stratigraphic characteristics. These elements are often distinguished as a result of a particular depositional process [PetroFAQ Glossary, 2022]. In the broadest sense, geobodies can refer to a stratigraphic subdivision such as channels, deltas, turbidite fans and salt domes, or structural elements such as fractures and faults.

1.2 Research motivations

The interpretation of seismic reflection data plays an essential role in understanding and characterising the subsurface. If 3D seismic reflection has been predominantly used in oil and gas exploration, its scope of applications extends to geothermal, subsurface energy storage, carbon capture and storage (CCS), mineral exploration, near-surface applications for engineering and environmental and monitoring subsurface change.

The interpretation of seismic data faces many sources of uncertainty and bias. First, the uncertainties related to the data quality inherent to the seismic acquisition and processing, such as ambient noise, artefacts, multiples, tuning effect, resolution, etc. Then, the uncertainties related to the interpretation process itself. Indeed, the structural and stratigraphic interpretation of seismic data can be highly subjective and open to interpreter bias. A single seismic image can lead to multiple geologically valid interpretations, especially in complex geological settings [Bond et al., 2007]. Interpretational uncertainty remains difficult to assess due to its subjective nature and significantly impacts subsurface modelling workflows and decision making.

Adequate interpretational uncertainty assessment requires multiple seismic interpretations. Manual interpretation of 3D seismic data is labour-intensive and is often influenced by the interpreter's prior knowledge. Therefore, an effort is made here to automate some seismic interpretation tasks to provide quick and reproducible interpretations.

Recent advances in machine learning have shown promising results in automating seismic interpretation for some tasks such as salt classification [Waldeland and Solberg, 2017] or fault detection [Wu et al., 2019]. Machine learning and deep learning are well adapted to seismic interpretation by their ability to capture hidden tendencies inside the data and to provide reproducible interpretation models. Moreover, machine learning models can generate multiple interpretational scenarios and thus help assess the uncertainties of the interpretations.

However, machine learning applications in seismic interpretation still face significant problems of generalisation and explainability. The first drawback is the lack of generalisation: even though some applications have demonstrated promising results when interpreting a single seismic survey, these models struggle to adapt to other seismic surveys with different geological settings, seismic resolutions, frequencies, noise levels, etc. [Lu, 2019, Zhu et al., 2020]. The other main drawback is the lack of explainability of the machine learning approaches: despite efforts toward more interpretability in machine learning models [Carvalho et al., 2019, Gilpin et al., 2019, Lundberg et al., 2020], extracting clear causality between input data and model predictions remains one of the most critical problems in machine learning today.

The primary motivation for this project is to combine machine learning approaches and geological domain knowledge to ensure meaningful geological information is extracted from 3D seismic reflection data. This research aims to use the advantages of machine learning to investigate aspects of 3D seismic data interpretation with fully data-driven workflows. More specifically, this thesis will tackle the problem of geobody detection and extraction from 3D seismic data.

1.3 Statement of the problem

1.3.1 Geobodies and their detection in seismic data

To find geobodies in seismic data, the primary attribute used is the amplitude of the seismic reflectivity data. Amplitude reflectivity is sensitive to variations in acoustic impedance and, consequently, to lithology, porosity and fluid content. In sedimentary formations, amplitude variations define sequence boundaries and changes in the depositional environment. As a result, it can be used to detail the shape and size of geological bodies and identify variations of lithology and fluid within them.

Detecting and delineating geobodies improves understanding of the subsurface and encourages including diversity in the conceptual models. Identifying the nature of a geobody and extracting it provides a geological understanding of the seismic data and its interpretation.

Geobody characterisation is fundamental to accurate geological modelling as geobodies can be included in the geological model and will have an impact on reservoir prediction and characterisation. Indeed, the fine-scale heterogeneity introduced by the geobodies may significantly affect the dynamic behaviour of a reservoir [Haldorsen and Lake, 1984, Jackson and Muggeridge, 2000, Novakovic et al., 2002].

This work aims to make a general approach to geobody detection and delineation, and we demonstrate our approach on two types of geobodies: **turbidite fans** and **sill intrusions**. To show the generalisation of the approach, we chose two types of features that are (a)

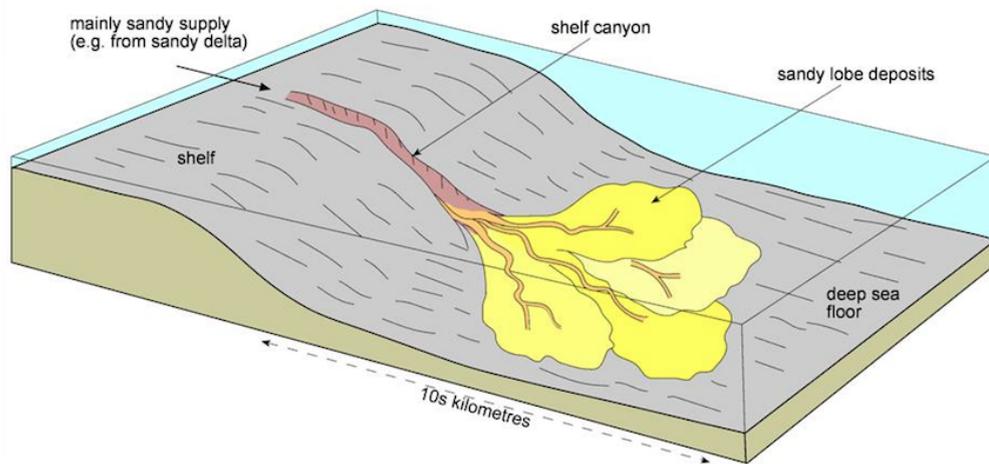


FIGURE 1.1: Facies model for a sand-rich submarine turbidite fan formed into lobe shapes that stretch out to the basin floor. From Nichols [2009].

different in terms of geometry and size and number and (b) geometrically complex and hard to detect as they are often subtle and time-consuming to find.

Turbidites are sediments deposited by a gravity flow along an underwater or sublacustrine slope (figure 1.1) forming clastic sedimentary rocks. Turbidite systems form excellent hydrocarbon reservoirs and traps that are important targets for oil and gas exploration. That is why their detection is a key issue. Turbidite reservoirs have complex heterogeneity and connectivity, so accurate characterisation is essential to understanding and predicting reservoir response. However, the characterisation of the turbidite fan is beyond the scope of this work; only their detection is studied here. In seismic data, turbidite fans are typically characterized by strong amplitude values and pinch-outs on the edges [Dorn, 1998]. Turbidite fans generally have a high acoustic impedance contrast compared to the encasing muds and shales [Bunt, 2015].

Sills are igneous magmatic intrusions that have intruded between older layers of sedimentary rocks. A magma flows and intrudes parent rock bodies following the path of the lowest pressure [Schofield et al., 2012] and according to the mechanical stress state imposed by the rock layering [Stephens et al., 2017, Walker et al., 2021]. Igneous intrusions are distinct in sills and dykes (figure 1.2), with sills generally described as concordant horizontal sheets following the geometry of the preexisting strata [Eldholm et al., 1995] and dykes as discordant intrusive inclined sheets that cut across the older

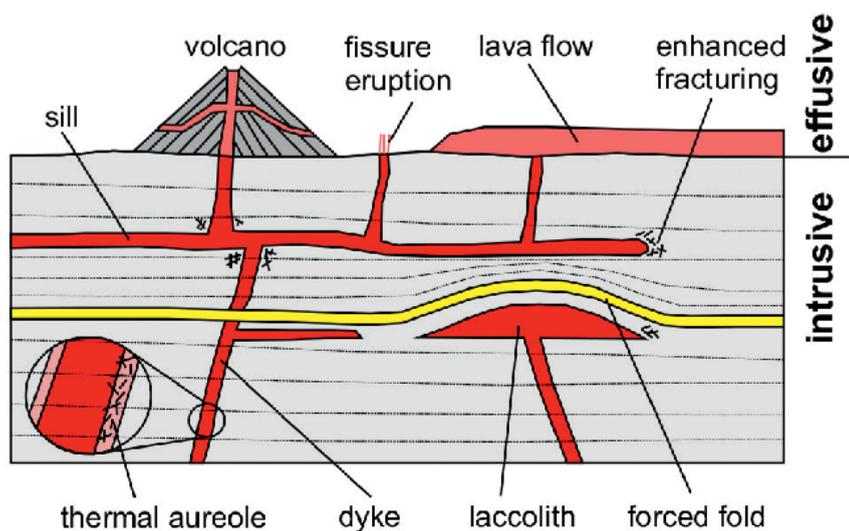


FIGURE 1.2: Types of magmatic bodies. Sills are concordant sheet intrusions following the preexisting rock beds. From Senger et al. [2017]

rocks, but it depends on the geometry of the strata. Sills and dykes are in close spatial association, often forming a connected intrusion network. The detection of sills and the analysis of their structures provide information about the basin evolution and the role they play in active hydrocarbon systems [Løtveit et al., 2013, Stephens et al., 2017]. Indeed, these intrusions are a source of heat and maturation for hydrocarbons and commonly change the petrophysical properties in the surrounding sediments [Planke et al., 2012]. They also form interconnected low-permeability zones, which can compartmentalise significant volumes of source and reservoir rock or serve as seal rocks [Rateau et al., 2013, Schofield et al., 2012]. Igneous intrusions have specific seismic signatures characterised by a high amplitude reflection due to high impedance when compared with the host sediments. Igneous intrusions can primarily be identified on seismic sections due to their tendency to cross-cut stratigraphy and their complex geometry [Thomson and Hutton, 2003, Thomson and Schofield, 2008]. Sills typically have a concave-upwards morphology [Du Toit, 1920], which produces the saucer-shaped geometries that help distinguish sills on seismic sections (figure 1.3).

The current geobody interpretation relies on the initial visual detection of a specific geobody by the interpreter, followed by a delineation of this geobody to extract it. The geologist detects the geobody directly on the seismic sections or the previously extracted horizons. This approach requires the interpreter to screen through the seismic

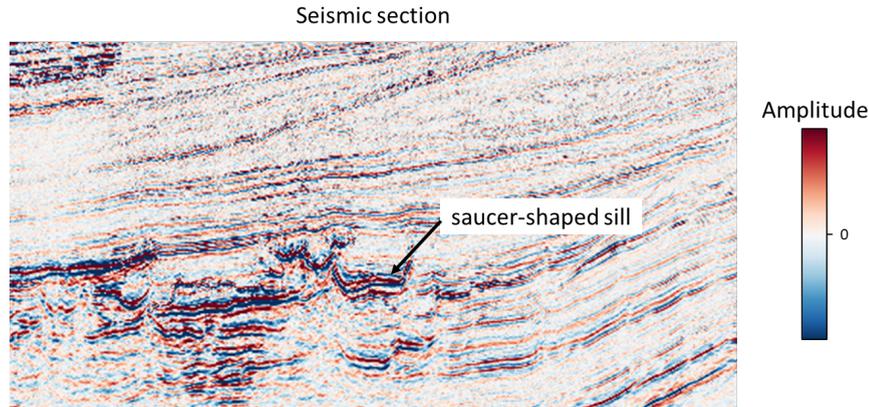


FIGURE 1.3: Saucer-shaped sill geometry on a seismic section (image from the Falkland Plateau Basin)

cube or horizon stack, looking for seismic anomalies based on their seismic attribute distribution and shape. However, interpreting geological features in seismic data is subjective and depends on visual perception. For instance, Froner et al. [2013] have shown that colour can bias the interpretation of geoscientific information and contributes to interpretation uncertainty and risk. Moreover, when a geobody is not following the main stratigraphy, the interpreter will not be able to characterise it appropriately with horizon-based interpretations. Once a geobody is detected, the geometry of the geobody is defined more finely by manipulating seismic attributes. We distinguish two main common approaches to geobody extraction: *manual delineation* and *clustering on seismic attributes*. In manual delineation, the interpreter either brushes over the geological features on the horizon or uses seed points to track and control the growth of the geobody in 3D. With clustering methods, a clustering algorithm groups pixels with similar seismic attribute responses into classes and extracts the class corresponding to the geobody detected [Schmidt et al., 2013].

1.3.2 Challenges

When working on automated geobody detection from 3D seismic data, we face two types of challenges: (i) the challenges inherent to the geobody detection task and (ii) the additional challenges consequent to automation.

1.3.2.1 Challenges in geobody detection

The first challenge in geobody detection results from the issues common to all seismic interpretation tasks: seismic reflection is an indirect and incomplete image of the subsurface from which lots of uncertainty arise. Seismic will have a varying signal-to-noise ratio, varying resolutions, varying reflector continuity, varying fault definition and possible processing artefacts. A geobody might be only incompletely imaged, missing parts of its body. Therefore, any geobody detection method must be robust to noise and adaptive to seismic resolution and quality.

The indirect and incomplete nature of seismic data also leads to interpretational uncertainty challenges. From the same seismic image, multiple interpretations can be made, all of which are possible. The uncertainties associated with seismic interpretation are complex to quantify. It is often necessary to incorporate a deep geological understanding to extract specific valid scenarios that the data allow.

Each type of geobody presents its set of challenges. For instance, not all geobodies can be characterised from 2D sections. In our turbidite fan example, the complete 3D shape is essential for an interpreter to recognise a fan confidently. Turbidite fans are often depicted as isolated, laterally extensive features, making their detection and delineation in 3D seismic data complex without prior knowledge of their position. On a 2D section, a turbidite fan can easily be mistaken for any other horizon. To illustrate this challenge, Figure 1.4 shows two sections crossing through two turbidite fans, but it is not straightforward to identify them from the section views. We can argue about high amplitudes or pinch-outs on the edges. Still, there is nothing immediately noticeable and exclusive to turbidite fans, making them simple to overlook during the interpretation process. A systematic interpretation of horizons and amplitude extraction is necessary to identify the fans with their characteristic shapes and amplitude distributions. Current geobody detection workflows do not include the 3D geometry properties of the object in a detection stage. Moreover, characterising a fan remains challenging even when working on the 3D shape of the object. The geometries of the fans vary depending on predominant grain size (figure 1.5), slope, basin floor angle, sediment flux, water depth and the degree of confinement of the turbidite system [Dodd et al., 2019]. If we add that

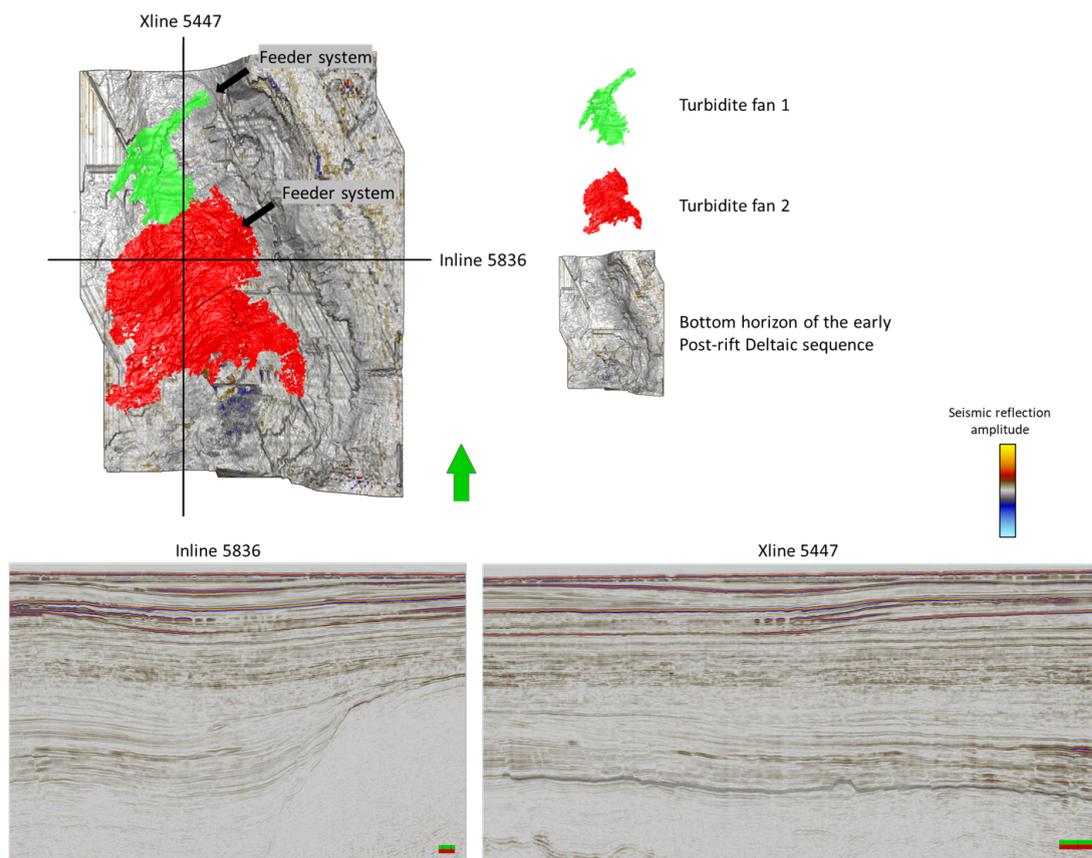


FIGURE 1.4: Section views in the seismic cube of two turbidite fans. This illustrates the complexity of detecting turbidite fans with only 2D information. Images from the North Falkland Basin.

the fan could be partially faulted or incomplete due to missing parts not imaged in the seismic, the range of shapes expands even further.

The main challenge with the interpretation of sills comes from the complexity of following sills laterally in the seismic data. Sill detection from 2D sections is less challenging than turbidite fan detection. Indeed, the bright reflections and characteristic shapes of sills (e.g. saucer-shaped) make them easier to spot when the section crosses the sill in its middle. However, following sills over inlines and crosslines must be carried out cautiously to avoid miss-ties. In particular, following the dykes that link sills together is difficult while necessary to understand the intrusion network. Seed-based auto-trackers would not do a good job extracting the entire intrusion network due to their predominant lateral search. Furthermore, many interpretation software packages do not allow the use of multiple z-horizons or structures, which are not adapted to the overlapping nature of intrusion networks. In addition, the automatic detection of sills will face the challenge

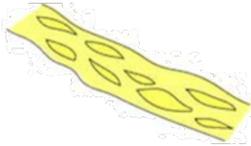
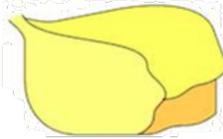
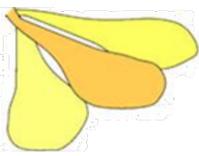
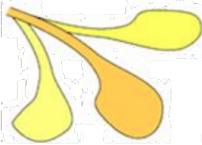
	Channels	Lobes
Sand-rich systems	braided channels 	channelised lobes 
Mixed sand and mud	channels and levee complexes 	depositional lobes 
Muddy systems	channels and levee complexes 	depositional lobes 

FIGURE 1.5: Schematic diagram showing the fan geometries that result from various sediment types. Adapted from Nichols [2009]

of the variety of sill shapes and sizes, making it difficult to extract logical rules for the classification of sills versus any other type of object in the seismic image.

1.3.2.2 Challenges in automation

The first challenges with automating a seismic interpretation task are flexibility, adaptability and robustness. A seismic image will vary greatly depending on the resolution, or the processing applied. Therefore, any automatic solution should be flexible and not hard-code any resolution-dependent rules. The subsurface presents a high diversity of geological structures and environments, and the solution should adapt. Finally, noise is a common issue in seismics. Thus, the solution should be robust to noise.

The second type of challenge is about communicating explainable and interpretable outcomes with some level of control over them. Indeed, automatic solutions and, more precisely, machine learning algorithms are sometimes defined as *black box* [Petch et al., 2022] since it can be challenging to comprehend how they operate once they have been

trained. This lack of interpretability in the outcomes of the solution can undermine trust in those solutions. An automatic solution to geobody detection should let the interpreter see and control how the solution works. Interactive tools at every step of the geobody detection workflow could help users understand why they receive a specific outcome and adapt it where necessary.

The third challenge is the quality control (QC) of the solution. QC is an essential aspect of automated solutions where results get validated, but the QC of automated results should not be more laborious than the manual interpretation itself. The quality of the outcomes must be measurable or simple to assess.

1.4 Research methodology

The research methodology to develop and assess geobody detection methods follows the standard practice in data science and machine learning with a clear distinction between a *validation set* and a *test set*. A *validation set* is used to understand the problem, characterise it and evaluate the methods during the development. The methods are selected and first assessed on this set of data. In contrast, the *test set* is completely unseen and unknown during the development phase to ensure that no bias is introduced toward a better, more informed result.

Since we are using an unsupervised approach to tackle our problem, no *training set* is necessary to train our models compared to supervised learning methods. The choice of the developed unsupervised methodology was dictated by the constraints of the problem and the data. These constraints and choices are explained and justified in more detail in Chapters 3 and 4. The main advantage of the unsupervised approach is that no training of the models is necessary. Therefore, no bias is introduced by a labelled training database, which makes the methods easier to adapt to other seismic data or problems.

To carry out this research, we have at our disposal a significant size, high-quality dataset from the North Falkland Basin (NFB) (figure 1.6 a & b), provided by the Falkland Island Government. This dataset contains a modern 3D seismic cube, processed on

45000 km^2 . It was selected for: (i) the diversity and quality of the data available with modern, high-resolution 3D seismic reflectivity data as well as logs and core data of a comprehensive nature; (ii) being heavily studied and characterised in recent years; (iii) displaying several Cretaceous-aged fan bodies [Bunt, 2015, Richards et al., 2006, 1996]. The turbidite fans have been manually interpreted in the 3D seismic images and characterised with precise core analysis, wireline logs, geochemical analysis and other diverse data types [Dodd et al., 2019]. Core data have been extracted, sampling several locations of the fans and proving the sand bodies to be oil and gas reservoirs [MacAulay, 2015].

The North Falkland seismic cube was cropped to cover the location of two known turbidite fans, which constitutes our *validation dataset*. The amplitude extraction maps were interpreted using Schlumberger's Petrel (figure 1.6.c) and with prior knowledge of the location of the two fans. This amplitude extraction will be used as ground truth for evaluating our turbidite fan detection and extraction method during the validation process presented throughout chapter 3 and chapter 4. The fan to the North is called the *North Sea Lion fan (NSL)*, and the one south to it is the *Sea Lion fan (SL)*. NSL is recognizable thanks to its long, well-preserved system feeder (figure 1.6.c). SL is recognizable with its double entry point system feeder. Variations of amplitude on the fan extraction maps represent depositional trends related to lithological variations and fluid content (hydrocarbon). The other advantage of working with only a cropped area during the development phase of the methods is making it easier and quicker to process. Indeed, the NFB seismic survey is a voluminous dataset (78GB), leading to memory and computation challenges.

We are using two *test datasets* and applying the geobody detection methods to them. Both tests are presented in chapter 5. The first test consists of detecting fans in the other areas of the NFB seismic survey. Thus, the idea is to create and validate methods on a sub-cube of the seismic image with two known fans and then apply them to the broader survey and test if we can retrieve the other fans present inside the survey. In the second test, the methods are applied to another seismic survey, the FISA 3D survey from the Falkland Plateau Basin (FPB) (figure 1.6 a), and another type of geobodies, sill intrusions. In this second test, the idea is to evaluate the robustness and the adaptability of the method to

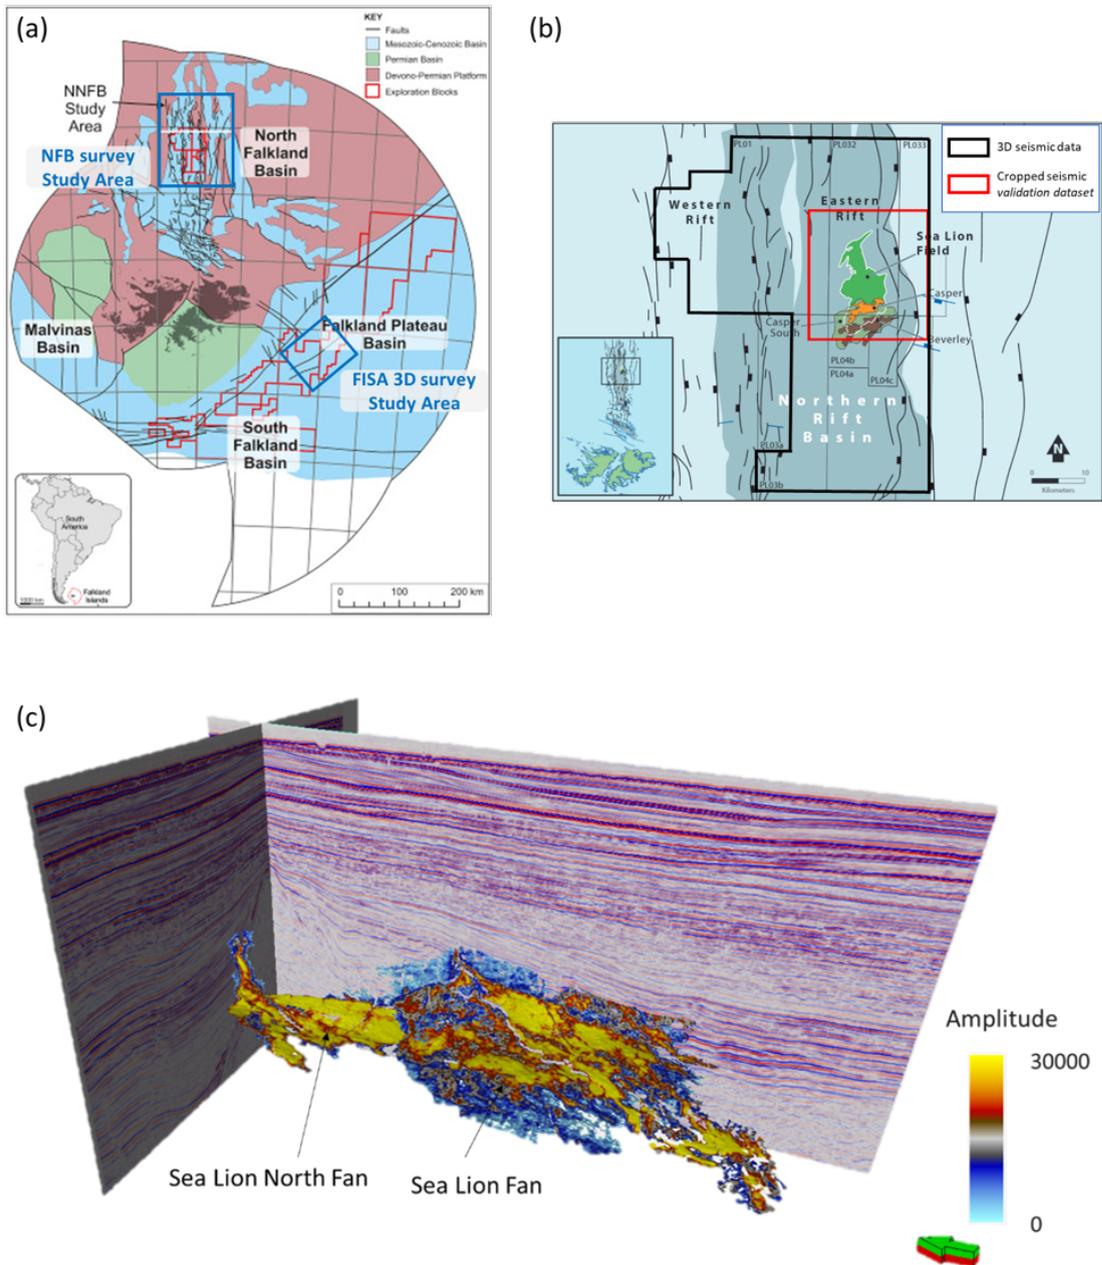


FIGURE 1.6: (a) Localisation of the different Falkland Basins. Modified from Jones et al. [2019]. (b) Localisation of the validation dataset, cropped from the North Falkland Basin. Modified from Bunt [2015]. (c) Sea Lion Fan and Sea lion North Fan amplitude extraction maps were manually detected and extracted using Schlumberger’s Petrel. Variations of amplitude on the fan extraction maps represent depositional trends related to lithological variations and fluid content (hydrocarbon).

a different geologic setting, seismic quality and type of geobodies. The basin has been affected by an Early Cretaceous-aged system of magmatic intrusions [Richards et al., 2013, Stanca et al., 2022]. The FISA 3D dataset is even more voluminous than the NFB survey (128GB). Working with real, basin-scaled seismic images enables testing for real complexity and demonstrates the viability of the proposed solutions in real-world case studies. Too often, research is being demonstrated on synthetic or small-scale studies, potentially leading to a lack of practicality for real case usage.

1.5 Thesis outline

The outline of this thesis is articulated around a novel workflow for geobody detection, which was developed for this work (figure 1.7). This workflow is subdivided into five steps but can roughly be summarised in two main phases. The first phase of seismic segmentation, which aims to isolate all objects present in the seismic, is presented in chapter 3. The second phase, which aims to find a specific type of object, is shown in chapter 4. Finally, chapter 5 presents the two application cases for the search of turbidite fans and sills in 3D seismic.

The thesis is divided into six chapters and structured as follows:

Chapter 2: Literature review – Automation in seismic interpretation

This chapter gives a brief overview of the recent evolution in seismic interpretation and the efforts toward more systematic and automatic seismic interpretation workflows. The chapter starts with a review of the main development in the interpretation of horizons, from first auto-trackers to global interpretation methods. In the second part, we present recent results of machine learning and deep learning solutions applied to seismic interpretation and the challenges that a representative training dataset poses. Then, we introduce works and research more specific to detecting and extracting geobodies from the seismic cube. Finally, the chapter closes with a presentation of the challenges of quantifying uncertainty in seismic interpretation and the advances made in this area.

Chapter 3: Full Seismic Segmentation with Point Cloud Representation and Density Clustering

This chapter introduces the methods for a novel seismic segmentation solution developed for this thesis. The seismic segmentation allows a division of the seismic volume into small objects at the scale of the geological structure. The segmentation constitutes the first phase of the geobody detection workflow. The solution relies on a sparse representation of the seismic as a point cloud. This point cloud is then clustered based on the spatial density-connectivity of the points in an unsupervised fashion. Methods and metrics for quality control (QC) of the segmentation results and optimisation of the clustering are also presented. Illustrations and validations of the different methods presented are computed on the *validation dataset*.

Chapter 4: Detection of 3D geobodies

This chapter presents the second part of the methods, now focusing on geobody detection from a segmented seismic. Segmentation results produced by the methods introduced in chapter 3 form a reduced-order representation that enables working on the set of individual seismic objects directly. We introduce a set of features used to characterise the 3D objects based on their geometry and amplitude distribution. Finally, we present a closest object retrieval approach that allows retrieving n (a fixed number) objects most similar to an example input by the interpreter. Illustrations and validations of the different methods presented are computed on the *validation dataset*.

Chapter 5: Geobody detection in 3D seismic - Application to real case studies

In this chapter, the geobody detection methods introduced in chapter 3 and chapter 4 are applied to two case studies, the *test datasets*. The first one, from the North Falkland Basin, applies the workflow to fan detection. Then, to further test the method, we apply it to a dataset completely unseen at the time of the development, the Falkland Plateau Basin, and a different type of geobodies: sill intrusions.

Chapter 6: Summary, Discussion and Conclusion

This chapter summarizes the main findings and conclusions presented in this thesis and compares and contrasts them to the thesis's initial objectives. Additionally, it offers reflections and suggestions for future development and application works, necessary improvements and unanswered questions.

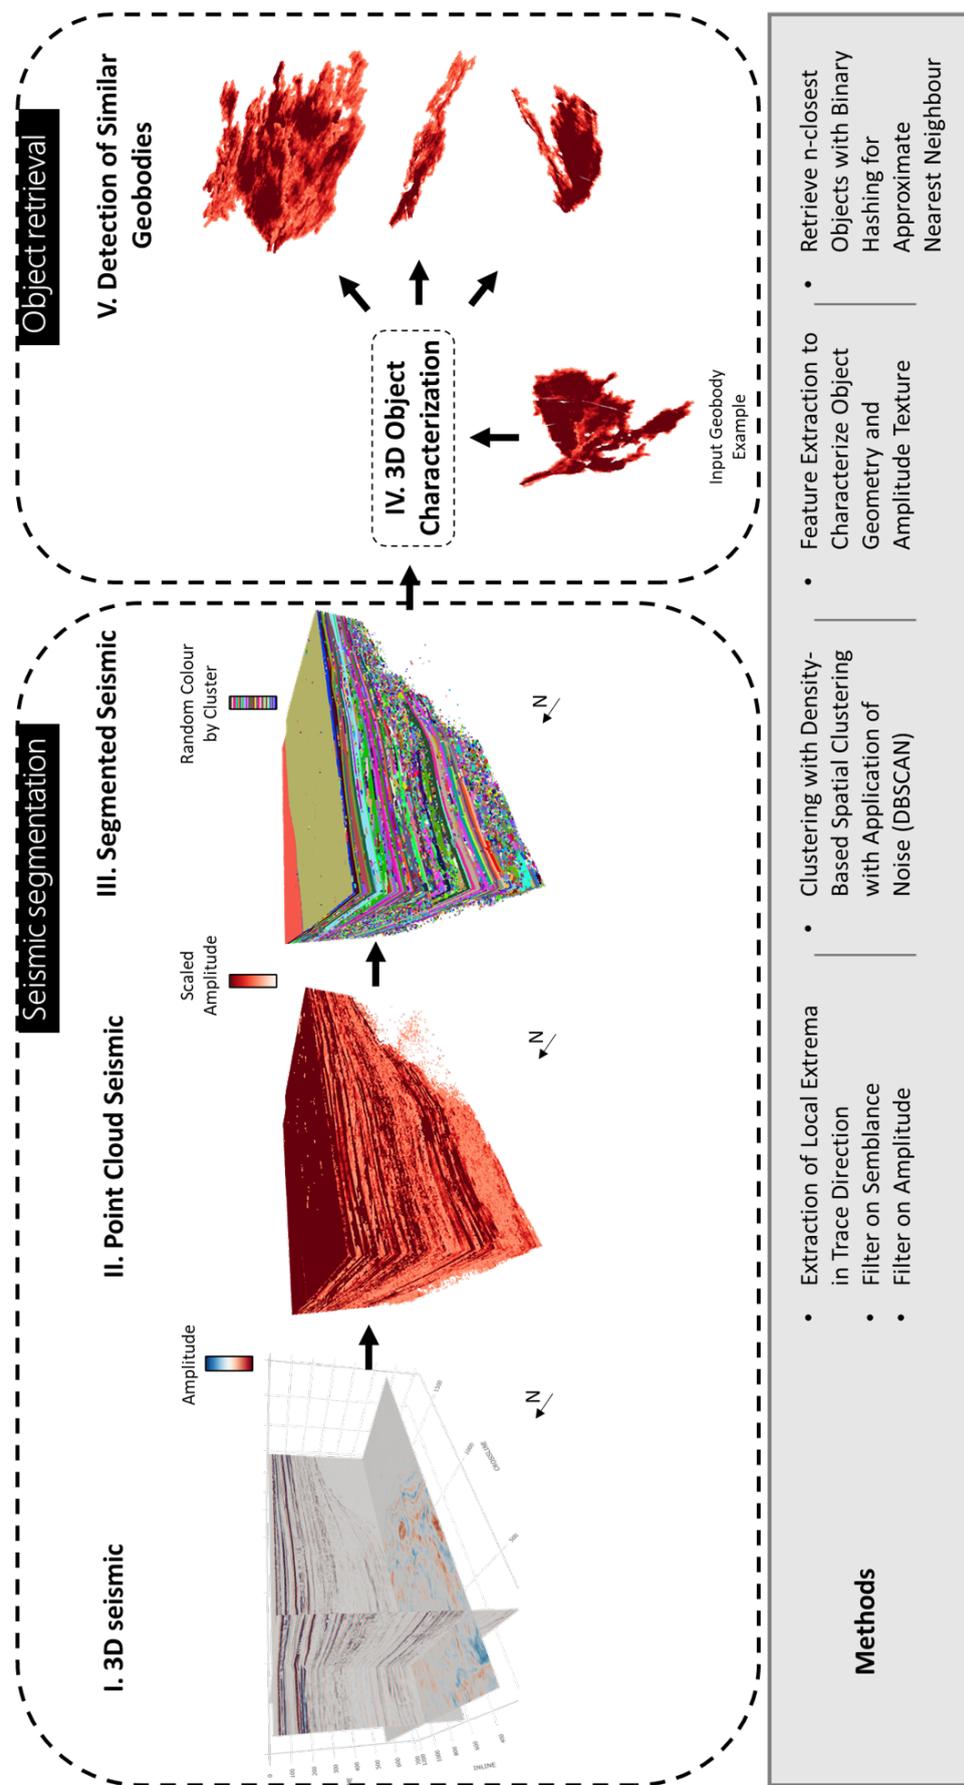


FIGURE 1.7: The geobody detection workflow proposed in this thesis. On top is a schematic representation of the workflow with a 3D seismic as input and potential geobodies as outputs. At the bottom are the main highlights of the methods used at every step of the workflow. The methods regarding seismic segmentation are introduced in chapter 3, and the methods regarding object retrieval in chapter 4.

Chapter 2

Literature review – Automation in seismic interpretation

Since the first 3D seismic survey shot in 1967 [Dorn, 1998], seismic interpretation has greatly evolved towards more automation, helping extract more information, reduce interpreter bias, and speed up performance. This chapter gives a brief overview of the recent evolution in seismic interpretation, starting with automation for horizon interpretation -a main focus in seismic interpretation- with the presentation of auto-trackers and global interpretation methods. In the second part, we present a new generation of methods, machine learning and deep learning techniques that transform the approach to automation in seismic interpretation applications. Next, we introduce more specifically the current techniques and research on the subject of geobody extraction in a seismic volume. Finally, this chapter closes with a presentation of research on the quantification of uncertainty in seismic interpretation, a complex issue to consider when automating a task.

2.1 Interpretation of horizons from 3D seismic data

2.1.1 Picking horizons

Traditional 3D seismic structural interpretation workflow produces horizon maps depicting the main subsurface structures. Traditionally, the interpreter manually extracts the horizon by manual picking. In manual picking, the horizons are interpreted by picking a reflector on 2D slices and loop tying over a regular grid on crosslines and inlines [Bacon et al., 2007, Badley, 1985]. To select the event to be picked, the interpreter looks for some local continuity in the data and local wavelet similarity. However, this manual picking approach is time-consuming and requires regular quality controls to avoid miss-ties.

2.1.2 Horizon auto-trackers

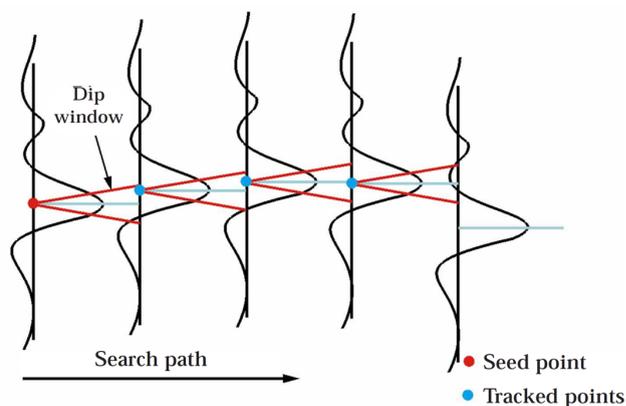


FIGURE 2.1: Sketch of a simple auto-picking algorithm. Figure from Dorn [1998].

With technological innovation, computer-assisted horizon trackers have started appearing, and significant efforts have been made to automate seismic interpretation. Howard [1991] first proposed a solution for auto-tracking (or auto-picking) horizons, expending interpretation away from a seed point. Data points are tracked over tiles based on an acceptance criterion considering amplitude values of peaks, troughs and zero-crossings. Gradually, the interpreter inputs seed points until a complete horizon is created. Since then, auto-trackers have been widely used, and derivatives of that initial implementation

have emerged. All these derivative implementations have improved the accuracy and performance of auto-trackers. Still, they are following the same general principle of seismic events tracked away from an input seed point by correlation of local waveforms between neighbouring seismic traces [Dorn, 1998] (figure 2.1).

In general, auto-pickers perform poorly when tracking across faults or when applied to complex geology with amplitude response distortion. Indeed, auto-pickers are based on local trace similarity and therefore are sensitive to variations in signal-to-noise ratio in the data and assume that the data are locally continuous, smooth, and consistent. The weakness of these seed-based auto-trackers is that they use only a small fraction of the data at a time and thus miss a more optimal global interpretation [Hoyes and Cheret, 2011]. To overcome this weakness, a new generation of seismic interpretation techniques, referred to as global interpretation methods, has emerged over the last decades with the improvement of computing performance.

2.1.3 Global seismic interpretation methods

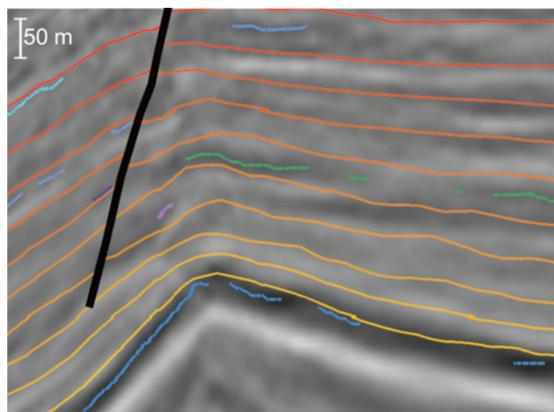


FIGURE 2.2: Systematic horizon extraction from global interpretation methods. The discontinuous cool-colored lines are the auto-tracked horizon patches interpreted by a global interpretation method. The continuous warm-colored lines represent the relative geological time interpolated between the patches. Figure from Labrunye and Carn [2015].

Global (or Full-volume) seismic interpretation methods are a new generation of seismic interpretation tools that produce a volumetric interpretation of a full seismic data volume. Instead of manually picking horizons, these methods systematically extract high-density

of horizons on the entire seismic cube, thus generating relative geologic time (RGT) volumes [Stark, 2004] (figure 2.2).

Hoyes and Cheret [2011] reviewed the different global interpretation methods and classified them into three categories. (i) Dip-driven methods [De Groot et al., 2010, Lomask et al., 2006] use local dip and azimuth information at each grid position within the volume. All seismic events are tracked by fitting horizons to the local dip calculated from the gradient of the seismic data at every sample position. (ii) Horizon patches methods [Borgos et al., 2003, 2005] groups the seismic signals along reflectors into small surfaces based on similarities in the shape of the seismic wavelet around the reflectors. The shape of each reflector can be represented through a set of attributes with one-point support, selected at traces extrema positions. And the horizon patches can be merged in a second step to form larger horizons. (iii) Global optimisation methods [Lacaze et al., 2017, Pauget et al., 2009a, Paumard et al., 2019, Possee et al., 2022] track simultaneously horizons by creating links between elementary seismic horizon patches over a coarse seismic grid (figure 2.3). The correlation between seismic samples is detected, minimising a cost function over the entire grid. The cost function relates to the seismic similarities of the seismic signals and the distances between the seismic samples.

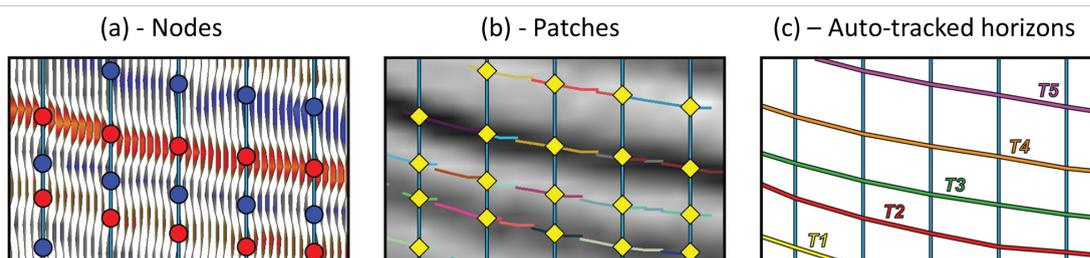


FIGURE 2.3: Workflow computing RGT model by establishing links among elementary horizon patches. The link between patches is created by minimizing a cost function based on the signal amplitude of neighbouring traces. Figure from Paumard et al. [2019].

Global interpretation solutions increase the amount of geologic information that is routinely extracted from the seismic image [Gogia et al., 2020]. These seismic interpretation solutions are now available in several commercial software packages and allow

the interpreter to perform high-resolution interpretation of 3D seismic data in a timely manner.

These advanced seismic interpretation methods highlight the importance of trace-to-trace similarities between seismic signals to track seismic events. They also highlight the importance of working on coarse grids or from samples from the seismic grid to get efficient automated solutions. Therefore, these two elements must be considered when developing a solution for an automatic geobody detection problem.

2.2 Machine learning and deep learning in seismic interpretation

Over the last decade, developments in machine learning and deep learning from outside the geoscience domain have led the way by which we interact with data and the way we approach automation to evolve [Larsen et al., 2018]. Many research works have emerged applying ML-based methods to geosciences. We first find applications to 1D data problems (e.g. well log classifications), and soon after, with the increasing popularity of convolutional neural network (CNN) [LeCun, 1989, LeCun and Bengio, 1995, Lecun et al., 2015] for image recognition and applications to seismic interpretation problems (2D and 3D).

2.2.1 Definitions

Machine learning distinguishes itself from other approaches by the use of statistical algorithms to “learn” information directly from data without relying on a predetermined equation or logical rules. These approaches are categorised as data-driven in opposition to conventional model-driven approaches. Machine learning methods are divided into supervised and unsupervised learning. Supervised methods rely on learning from labelled examples to apply to some new data where the labels are unknown. On the other hand, unsupervised methods do not require the data to be labelled. Instead, the

algorithm finds unknown patterns inside the data and groups them into clusters. Deep learning is a set of machine learning methods based on neural network models composed of many layers of neurons (deep).

2.2.2 Machine learning applied to seismic interpretation

Both supervised and unsupervised techniques have demonstrated interesting results for seismic interpretation. Strecker and Uden [2002] used Self-Organising Maps, an unsupervised approach, to cluster seismic over multiple post-stack seismic attributes (wavelet envelope, wavelet bandwidth, similarity). The results highlighted the ability of machine learning to embed multiple data features simultaneously to extract additional information from the combination of them (figure 2.4.a). Supervised learning enables the identification of specified seismic patterns like salt bodies, faults or seismic facies (figure 2.4.b). A common workflow is to treat seismic data as images and perform classification over small samples (patches). The model learns from some labelled seismic lines or blocks and applies it to others [Wrona et al., 2018].

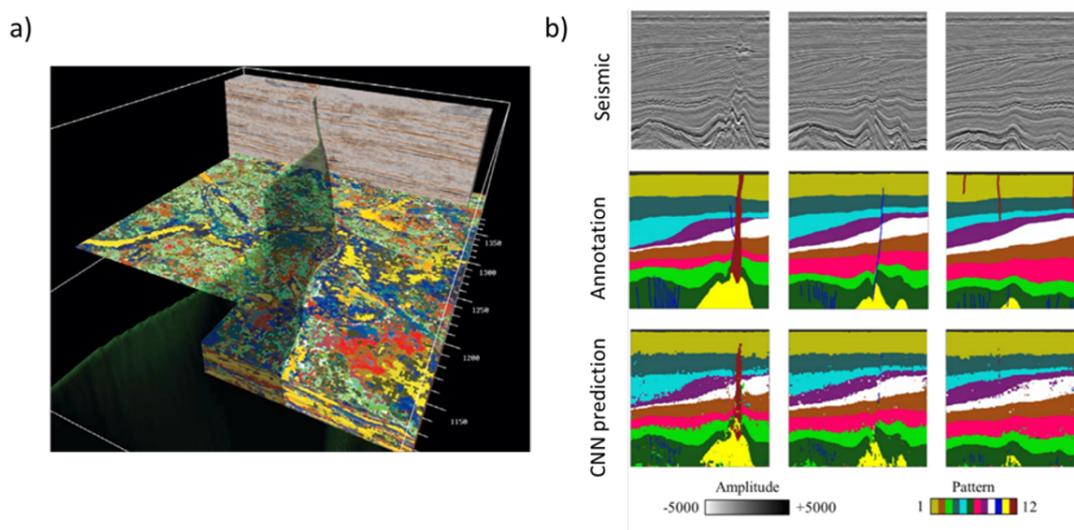


FIGURE 2.4: Unsupervised and supervised machine learning applied to seismic interpretation. (a) Unsupervised machine learning for clustering seismic patterns on a seismic cube given multiple seismic attributes. Clustering of the seismic cube reveals channel structures. Figure from Strecker and Uden [2002]. (b) Supervised machine learning, CNN, for seismic patterns classification, training from 2D line annotations and predicting labels on other lines. Figure modified from Di [2018].

2.2.3 Deep learning applied to seismic interpretation

Deep learning broadens the perspective on what could be extracted from the data for computer vision problems and hence for seismic interpretation. In particular, CNN displays impressive results for spatial feature extraction and image classification. For instance, Waldeland et al. [2018] applied CNN to classify salt in seismic images. The network is trained from a few seismic lines manually labelled and applied to every other line of the dataset. The classification is performed using a patch-wise method: the seismic lines are divided into small patches, and every patch is binary classified to either salt or “else”. One limitation of the patch-wise method is the necessity to be able to identify the seismic facies from a small patch. To overcome this limitation, a solution is to input the entire image and train a Neural Network to predict masks over the entire image. This task is called semantic segmentation and is solved using fully convolutional networks (FCN) [Long et al., 2015] or U-net architecture [Ronneberger et al., 2015]. For instance, Chevitaese et al. [2018] and Civitaese et al. [2019] applied a FCN to segment seismic images into different seismic facies, demonstrating results close to a human interpreter. These approaches demonstrated the use of deep learning for propagating manually-interpreted 2D slices to the 3D seismic cube. However, the interpretation is limited to the dataset on which the neural network has been trained. It needs a new labelling and training process every time it is applied to a new dataset. Another problem working only with 2D slices is that 2D information does not capture possible anisotropy in the third dimension. Hence, slice-to-slice segmentation does not get the 3D context necessary for accurate segmentation and semantic annotation. 3D U-Net architecture has been proposed to perform 3D semantic segmentation [Cicek et al., 2016], solving the anisotropy problem. With semantic segmentation, the network is trained to recognise patterns of stack reflectors rather than individualize seismic features. To interpret and delineate an individual structure such as a geobody, the segmentation method to be applied is instance segmentation. Instance segmentation refers to the detection and segmentation of objects into separated instances. The current state-of-the-art instance segmentation on images is performed using Mask R-CNN [He et al., 2017].

The problem of geobody detection can be framed as an instance segmentation problem,

where each geobody is an instance. But the application of instance segmentation to seismic interpretation faces two main challenges: (i) the adaptation to mask on 3D objects and, more importantly, (ii) the lack of a properly labelled seismic dataset that would satisfy the instance segmentation task.

2.2.4 Training dataset for seismic interpretation

2.2.4.1 Challenges with building a representative seismic dataset

The lack of a representative training dataset has been a recurring problem for machine learning and deep learning applications to seismic interpretations. First of all, part of the problem stems from the fact that there are no boundaries for a dataset representation of seismic examples: one cannot say that a labelled seismic database covers all case scenarios. Unlike MRI scanner applications -from which many deep learning applications originate-, where it is more or less possible to estimate the possible brain image scenarios, there is no such range of scenarios with the structures present in the subsurface. Secondly, the additional tricky aspect of seismic data is that a sample outside the training representation involves differences from the training representation not only in terms of structure or geographical provenance but also differences from the training representation in terms of processing. Finally, the labelling of data is specific to each case and type of seismic interpretation. The same seismic data can be labelled in very different ways depending on the level of interpretation -structural, stratigraphic- or the type of approach -patch classification, bounding boxes for object detection or masks semantic segmentation (figure 2.5). And if we take into account the uncertainties inherent to seismic interpretation, this would lead to having to create several labels for the same seismic event to integrate all the valid scenarios that the data allow. And of course, one must also be aware of the bias introduced by the interpreter when labelling the data.

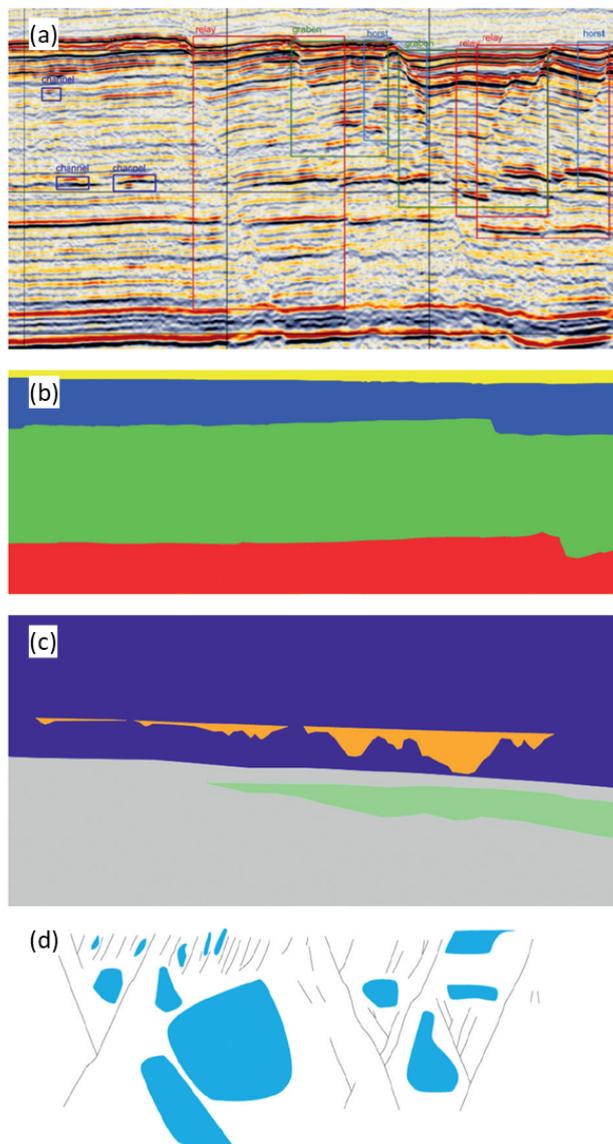


FIGURE 2.5: Different ways to label seismic data for machine learning tasks; (a) bounding box labels for subsurface object detection, (b) dense labels over an entire inline for stratigraphic unit segmentation, (c) dense labels for geobody detection and delineation, (d) sparse label sets on an inline section for use in automated fault interpretation. Figure from Larsen et al. [2018].

2.2.4.2 Open-source data initiatives

That said, there are more and more open data and open source dataset initiatives such as the Volve dataset or Netherlands off-shore F3 block seismic data [Birnie, 2021, NLOG, 2020]. The growing availability of three-dimensional seismic data to researchers will help alleviate these data problems and also provide benchmark data to compare the effectiveness of different approaches and methods. But still, geosciences have some

progress to make to catch up with the availability of databases offered in other industries using deep learning in terms of richness and quantity of labelled data.

2.2.4.3 Data augmentation and transfer learning

Data augmentation techniques allow for artificially increasing the number and diversity of samples in a database. These techniques are helpful in cases where relatively little data is available, as in the case of seismic interpretation. The primary data augmentation techniques apply rotation, translation, different scale/zoom levels, or adding Gaussian noise to the seismic image.

Transfer learning is another technique commonly used in deep learning to improve the training of a deep neural network when little training data is available. The principle of transfer learning is to use a pre-trained network (trained on another task and database). In the first step of the training, only the final layer (output) of the neural network is trained, and the neural weights of the other layers are frozen. Thus these neural weights, which have already proven effective, serve as a fixed feature extractor. Transfer learning reduces the amount of training data and training time required to achieve good accuracy [Tan et al., 2018]. The use of a pre-trained CNN has shown promising results in adapting to seismic texture identification [Dramschi and Lüthje, 2018].

2.2.4.4 Synthetic seismic data

Another way to solve the problem of insufficient or missing labelled data is to use synthetic data. The idea is to use forward numerical modelling to simulate synthetic seismic data from a geological model. Thus the geological model is known and can be used as ground truth to label the data. Wu et al. [2019] suggests that synthetic training data sets with realistic simulated structures can effectively train CNNs for fault interpretation in 3D seismic data. The results depend on the realism of the synthetic seismic data and the diversity incorporated in the synthetic data set.

Building realistic geological models and synthetic seismic data is an important research topic. For instance, Wu et al. [2020] proposes a workflow to build realistic folding and

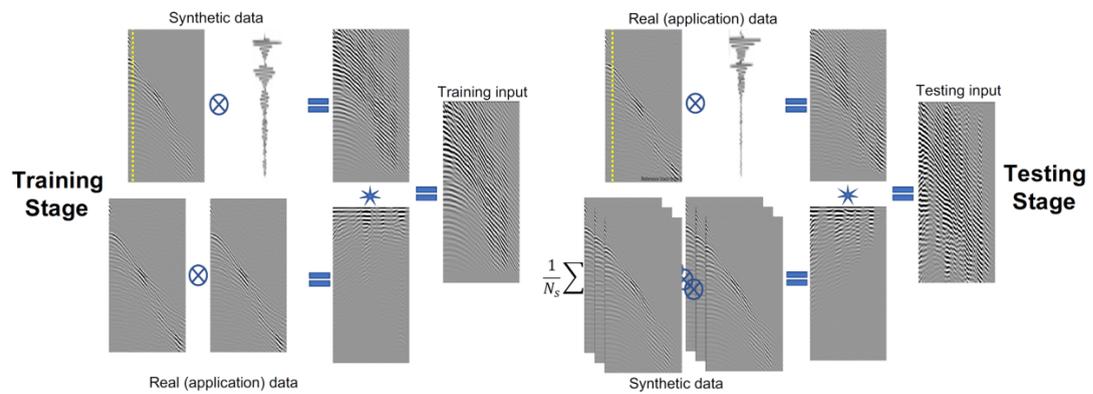


FIGURE 2.6: Workflow to incorporate real information in synthetic data at the training stage and synthetic information in real data at the inference stage. Left side: the proposed process used for producing the training data; Right side: the proposed process used for producing the testing/inference data. The idea is to use cross-correlation and convolution operations to mix real data with synthetic data so they have more similar distributions. Figure from Alkhalifah et al. [2021]

faulting features automatically and thus produce more realistic synthetic seismic that proved efficient to train a CNN for fault segmentation. However, in his conclusion, Wu et al. [2020] also mention that his workflow still needs to include geobodies and unconformities essential to create even more realistic structure models. These additional features could be added, but it is still complex to capture in synthetic models the level of realism and diversity of structures present in real data. Generative deep learning techniques such as generative Generative Adversarial Networks (GANs) [Goodfellow et al., 2014] or stable diffusion models [Rombach et al., 2022] also demonstrate promising perspectives to create realistic synthetic seismic data [Li et al., 2020]. To bridge the gap between synthetic and real data, Alkhalifah et al. [2021] suggests combining synthetic data with information from the real data at the training stage and information for the synthetic data used in training at the inference stage. Results using this preconditioned synthetic data helped improve the prediction for applications for microseismic event source location determination (figure 2.6).

2.3 State of the art on geobody extraction

In common seismic interpretation workflow, geobody interpretation is seen as a separate interpretation task performed after the main structural and stratigraphic interpretation. The structural and stratigraphic interpretation of the seismic data map the main horizons, faults and stratigraphic sequences, allowing a global understanding of the subsurface. The detection and delineation of geobodies produce a finer understanding of the constitutive of these stratigraphic sequences and help identify stratigraphic traps and heterogeneity that might impact the reservoir connectivity and migration pathways for fluid in the subsurface.

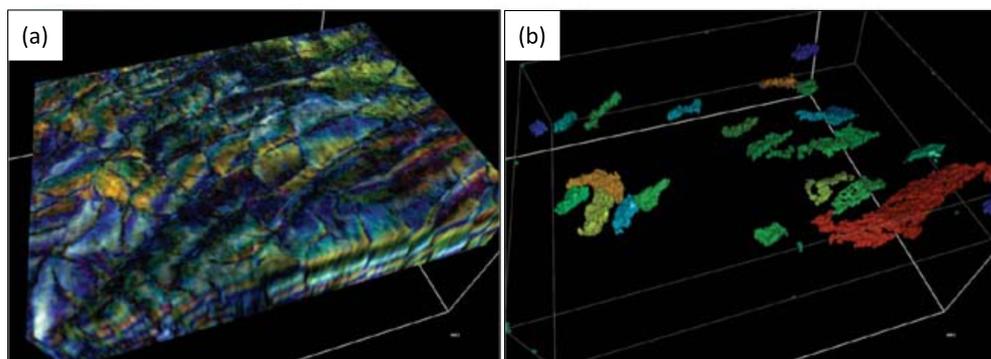


FIGURE 2.7: Body delineation from seismic attribute cube; (a) Spectral decomposition and RGB blend of the response at three different frequencies, and (b) bodies extracted using opacity technique. Figure from Henderson et al. [2008]

A common approach to geobody extraction is to display seismic attribute response on interpreted horizons [De Groot et al., 2016] and extract bodies using cut-off values on the seismic attribute. As described by Chopra and Marfurt [2005], a seismic attribute is a quantitative measure of a seismic characteristic of interest. These seismic characteristics can relate to the subsurface's geometrical and physical attributes. Seismic attributes can provide information on both the shape and size of geological bodies, as well as on the lithological and fluid variations within those bodies. Already in 1997, more than 60 frequent seismic properties were already catalogued by Chen and Sidney [1997] along with a description of their apparent significance and use. The most commonly used seismic attributes for geobody detection are spectral decomposition, coherency, root-mean-square amplitude (RMS), signal envelope and sweetness. The visualization of a seismic attribute on a surface allows the recognition of depositional patterns by

the interpreter. Then, cut-off values are applied to the attribute to extract the geobody [Henderson et al., 2008, 2007, Paton et al., 2011] (figure 2.7).

Another approach to extract geobodies from seismic attributes is to use edge detection algorithms such as Sobel filters or Canny edge detector [Canny, 1986, Kanopoulos et al., 1988] to extract geobody boundaries. For instance, Phillips and Fomel [2017] proposed a plane-wave Sobel attribute to efficiently create a sharper and more detailed image of channel boundaries.

There are also object-based geobody extraction methods where a rough model of the object input by the interpreter is deformed by optimization so that its shape matches the geobody information contained in the seismic data. For example, Ruiu et al. [2015] used nonuniform rational basis splines (NURBS) as a rough and editable model that can be easily drawn by the interpreter and optimized to extract channels from seismic images.

Machine learning and deep learning also have been applied to geobody extraction. For instance, Smith [2017] used unsupervised ML to illuminate particular classes of geobodies better. Le Bouteiller and Charl  y proposed a semi-supervised ML approach based on textural analysis to assist interpreters in identifying geobodies in 3D seismic data [Le Bouteiller and Charl  y, 2020, Le Bouteiller et al., 2019]. As for deep learning methods, Pham et al. [2019] suggested training a CNN on synthetic data to segment channels in 3D seismic cube. And de Groot et al. [2021] implemented Thalweg tracker, a combination of a tracker and a CNN to track a seismic event and extract a 3D geobody from a single seed position.

In this section, we mentioned several different approaches to geobody extraction and each of these approaches produce different outputs. Some methods are to help interpreters identify geobodies more easily, some are semi-automatic and only require minimal inputs (rough models or a minimal number of seed points), and others are fully automatic and produce the detection and delineation of the geobody. For instance, in Figure 2.8, we compare and contrast four typical outputs received from four different methods on the channel detection problem. The application and the advantages and disadvantages of each of these different geobody extraction methods are further discussed and put

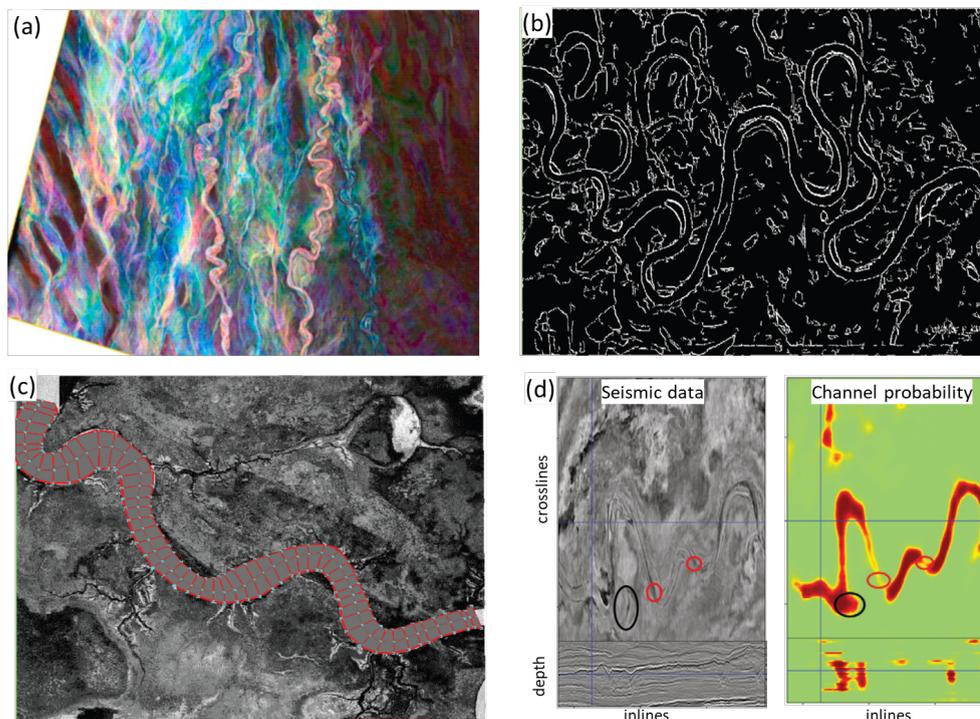


FIGURE 2.8: Comparison of the outputs obtained from four different methods to channel extraction; (a) seismic image enhancement with spectral decomposition (figure from Lawal et al. [2022]), (b) channel boundaries extraction with canny edge detector (figure from Ruiu et al. [2015]), (c) channel extraction with optimization of NURBS surfaces -object-based method- (figure from Ruiu et al. [2015]), and (d) channel detection and delineation with a CNN (figure from Pham et al. [2019]).

into perspective with our automatic geobody detection problem (fans and sills) in the introduction of Chapter 4.

2.4 Uncertainty in seismic interpretation

Another important aspect to be considered in any seismic interpretation workflow is the quantification of uncertainties. Uncertainties are inherent to the incomplete and indirect nature of seismic data. Bond et al. [2007] shown the presence of significant conceptual uncertainties in the interpretation of 2D seismic lines. A single seismic image could often support a variety of valid interpretations dependent on the geological understanding of the interpreter (figure 2.9). The same analysis has been extended to 3D seismic interpretation: Schaaf and Bond [2019] illustrated the uncertainties in

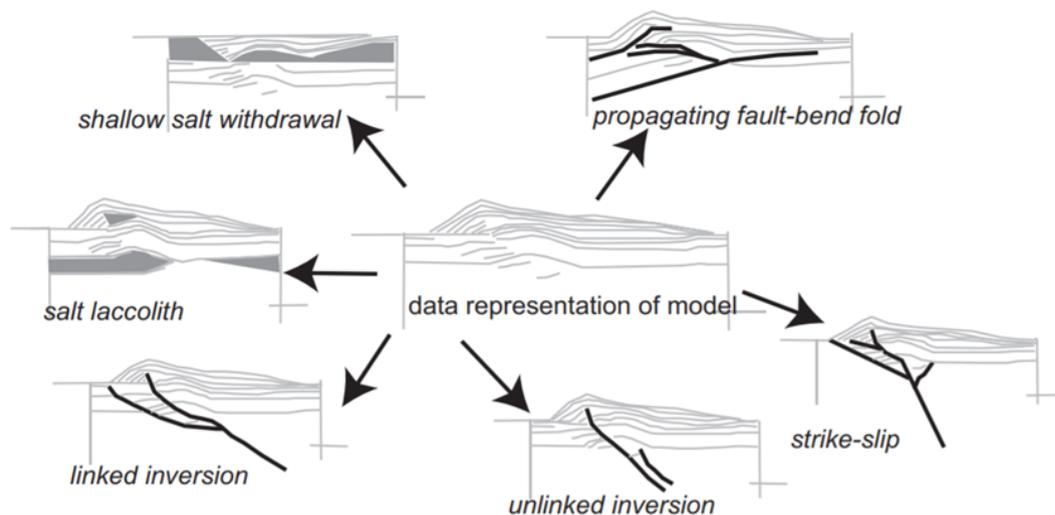


FIGURE 2.9: Examples of interpretations from a single 2D seismic image classified by tectonic settings. Figure from Bond et al. [2007].

3D seismic related to fault placement and fault topology. They emphasised the use of stochastic approaches, especially for the interpretation of areas of poor data quality.

Stochastic interpretation for structural geomodelling could help characterising the uncertainties by representing a range of possible interpretations. In addition, integrating more geological knowledge, in the form of logical rules, could ensure better exploration of the space of possible interpretations [Botter et al., 2016, Godefroy et al., 2021]. Given a range of possible scenarios, another problem is how to adequately rank the possible scenarios and compute their relative probability. To evaluate the likelihood of different structural interpretations from the same seismic image, Irakarama et al. [2019] proposed the use of misfit functions between synthetic seismic computed from the structural interpretations and the actual seismic data. The paper illustrated the need for conditions imposed on the data-space misfit function for reliable interpretation appraisal, conditions being dependent on the type of interpretation one wishes to undertake.

Some deep learning approaches also try to tackle uncertainty quantification in seismic interpretation. For instance, Bayesian neural networks allow us to learn a probability distribution and output predictions with uncertainty associated and have been applied to seismic facies segmentation and fault identification [Mosser et al., 2019, 2020].

2.5 Discussion

Significant advances in automation have been made to extract structural models (faults and horizons) from 3D seismic data. Different approaches are possible, all requiring more or less input from the interpreters and offering more or less control over the result, but these approaches all have in common the advantage of significantly reducing the labour-intensive work required to extract faults and horizons.

Machine learning, and in particular deep learning, has recently changed the field of seismic interpretation, offering the prospect of finer and more efficient automated interpretations by learning and integrating more geological knowledge by training neural networks. However, building representative training datasets for these neural networks to learn from has been a major concern, especially due to the lack of labelled seismic data.

Extraction of geobodies is often seen as a separate interpretation task not captured in initial horizon interpretation. Again, research has been carried out to move towards the automation of geobody extraction. However, progress must be made to provide automatic detection of the type of object without prior knowledge of the presence or position of the object. The 3D shape of the object and the amplitude distribution of the reflection are two primary criteria to allow confident identification of a geobody. We believe that using unsupervised or semi-supervised machine learning approaches can provide solutions generalisable to many types of geobodies and address the lack of training data.

Finally, uncertainty in seismic interpretation is a complex and complicated problem to integrate into an automation process and an active research subject. An important notion to consider is the notion of scenario-based interpretation, introduced by Bentley and Smith [2008], in opposition to anchor-based interpretation. Too often, a stochastic approach to the uncertainty problem leads to multiple interpretations, all being based on the same central (anchor) interpretation. In Bentley and Smith [2008], that notion is discussed for reservoir modelling, but the same applies to seismic interpretation. Seismic interpretation should aim to output multiple-deterministic scenarios that the

seismic data allow rather than minor variations of the same interpretation scenario, thus creating explicit dependencies between data uncertainty and the different possible geological concepts embedded in each scenario.

Chapter 3

Full Seismic Segmentation with Point Cloud Representation and Density Clustering

3.1 Introduction

3.1.1 Objectives and requirements

This chapter presents the methods used throughout the thesis to perform a Full Seismic Segmentation.

Segmentation is the division of an object into multiple segments, also known as regions or objects. The purpose of segmentation is to simplify the representation into something more meaningful and easier to analyse, gaining an understanding of the structures embedded in the data. In this thesis, by seismic segmentation, we mean the division of the seismic data into small objects at the scale of the geologic structure (figure 3.1.a). The output of the seismic segmentation is an ensemble of objects spatially delimited. Here, seismic segmentation has the ultimate purpose of compressing geological structural information embedded into the seismic data to a manageable size for interpretation without loss of meaningful interpretable information. In this chapter, we are not yet

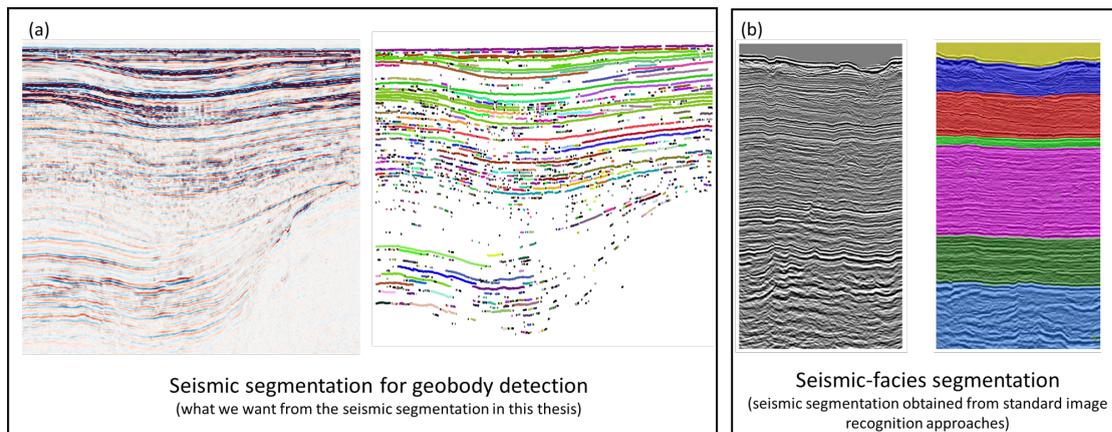


FIGURE 3.1: Comparison between the output needed from seismic segmentation for geobody detection (a) and from seismic segmentation at a seismic-facies scale (b) which people usually refer to when talking about seismic segmentation. Figure 3.1.b. adapted from Civitarese et al. [2019].

looking to label the objects with any semantics. Instead, the segmentation only labels the objects as a cluster or as some noise. We introduce the labelling of the cluster to meaningful geological structures in chapter 4.

Seismic segmentation at the scale of the geologic structure faces two main challenges: (i) working with objects of heterogeneous scales and (ii) working on big data volumes. If extracting large horizons is now a well-studied problem, segmenting at the same time subtle features such as turbidite fans and large features such as horizons remains very complex. Furthermore, segmentation on large datasets can be very computationally demanding, making it a hindrance for many approaches.

We present a novel method for global seismic interpretation which tackles segmentation as a spatial clustering problem, looking for *spatial continuity* between the seismic wavelet information and then optimising the groupings to uncover geologically interesting features. The two key advantages of our approach are (i) the ability to capture both large horizons and subtle objects and (ii) being computationally efficient, especially when dealing with large seismic volumes.

Seismic segmentation takes us from Raw 3D seismic data to a segmented volume of objects that potentially represent geobodies. Delimiting seismic data into objects offers the possibility of working on the geometric characteristics of these structures. To then detect geological features within the segmented volume, we must use the geometric

attributes of the segments. Thus, the segmentation of all objects in the seismic data is a prerequisite to object detection.

To be adapted to real data, the seismic segmentation method must be:

- (i) Performed at the scale of the seismic reflector, not at the scale of the stratigraphic units.
- (ii) Adaptive to the dimension of the objects to be detected. Indeed, the segmentation must capture features from large, laterally extensive horizons down to subtle depositional features, like channels or turbidite fans. The segmentation method should not assume continuity over the entire cube or between fault surfaces.
- (iii) Fast. When looking for small features in large seismic volumes, we are essentially looking for needles (e.g. channels or lobes) in a haystack (seismic cubes with billions of voxels). Segmentation groups the voxels by some similarity measure to reduce the problem complexity; we are now looking for a needle amongst fewer bundles of straw. Nevertheless, we need to ensure that the segmentation task itself does not take longer than just hunting through all the data. So, we are looking for a high-performance segmentation method to apply to sizeable seismic volumes, looking for subtle, hidden features.

3.1.2 Overview of existing methods

Full Volume or Global Interpretation methods have changed the way we interpret seismic data by extracting an ensemble of horizons from the entire cube to build relative geologic time volumes [Stark, 2004], replacing manual picking in many cases. As manual picking is slow, pain-staking work, Full Volume Interpretation provides considerable benefits in terms of time and cost. In addition, it increases the amount of useful geologic information extracted from the seismic data volume as a dense set of horizons gets created [Hoyes and Cheret, 2011].

There are three main approaches to the seismic segmentation problem; people usually state this problem as a region-growing problem, an optimisation problem, or an image recognition problem.

Most common approaches describe the interpretation of seismic horizons as a region-growing problem, using multi-horizon auto-tracking algorithms to track horizons away from bin samples [De Bruin and Bouanga, 2007, De Bruin et al., 2006, de Groot et al., 2010, Qayyum et al., 2012]. First, region-growing methods start with seed points, picked either manually or automatically. Then, they look at the surrounding pixels of the original seed points and decide if they should be included in the region. And they iterate upon this process until the region-grow is complete. However, the quality of the segmentation heavily relies on the number and location of the seed points. Besides, these implementations, based on iterative-process, are computationally-intensive [Poux et al., 2020].

Alternatively, we can treat interpretation as an optimisation problem, where we see to minimise the dissimilarity between seismic samples of the same sequence. The cost function is to minimise the distance and dissimilarity between the seismic samples belonging to the same seismic sequence. The optimisation process finds a global minimum of the cost function. Thus the optimisation produces an optimum global seismic model. However, these implementations are also computationally-intensive and often require solving the optimisation on a coarser grid [Pauget et al., 2009b, Paumard et al., 2019]. Working on a coarser grid leads to a suboptimal solution which might mean, for our application, missing subtle features.

Finally, other approaches present the seismic interpretation problem as an image recognition problem [Chevitarese et al., 2018, Waldeland and Solberg, 2017, Zhao, 2018]. Indeed, the recent development in the field of Deep Learning has led to a significant improvement in solving image recognition problems. The standard approach trains a deep neural network to associate a patch to a label. A patch is a small part of an image with a defined size. Thus, in seismic interpretation, the seismic gets divided into patches. Every patch of the seismic gets labelled by the neural network. Classifying the entire seismic cube into labels provides a segmentation of the seismic by the labels. However,

by doing this, the network is trained to recognise patterns of a stack of reflectors and seismic-facies (figure 3.1.b) rather than segment the seismic features at the object scale, which is a requirement for the segmentation we need here.

3.1.3 Our novel method

Our approach frames the seismic segmentation problem in a new way by considering it as a spatial clustering problem, making it fast to solve and independent of the size of the structures. This novel way of approaching the seismic interpretation problem is made possible by manipulating the seismic data as a sparse object. Instead of working with a cube full of data, we extract or filter a sparse representation that we store in a *point cloud* format [Corlay et al., 2020]. That sparse representation aims at creating spatial distance in-between the seismic features. How we convert the seismic to a point cloud is crucial as it will essentially lose information in favour of a simpler problem. Every point of the resulting point cloud is then spatially clustered. Each cluster represents a distinct seismic feature. We illustrate this approach in Figure 3.2.

Although this workflow is novel, it is interesting to note that a clustering approach to seismic interpretation is seen as very natural as it leads to the apparent interpreter's output - detect results and group (cluster) them into specific configurations or geological shapes (horizons, geobodies, etc.). Clustering merely mimics the manual interpreter's approach by aligning shape configurations together as a consistent group that represents a particular geological structure.

In this chapter, we first show how to extract a point cloud from seismic data; then, we segment the point cloud into distinct geological features. Next, we present how the results can be quality controlled and optimised from that seismic segmentation. We validate every method presented in this chapter by applying it to the well-studied *validation dataset* described in *Chapter 1 Research methodology*.

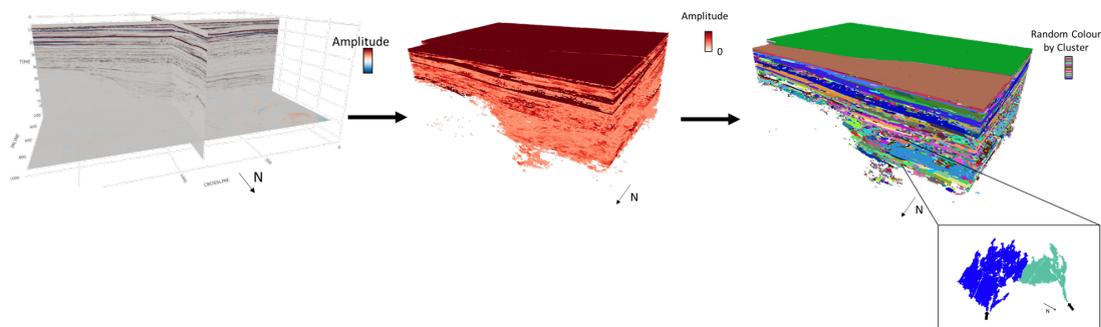


FIGURE 3.2: Workflow of the proposed seismic segmentation method. A point cloud gets extracted from the seismic. This point cloud is then segmented into individualized seismic features. The size of the segmented objects can range from extensive horizons covering the entire cube to more subtle features such as turbidite fans illustrated in the zoom here.

3.2 Point Cloud Seismic

3.2.1 Why working with a seismic data point cloud?

The information required for the interpretation may be extracted from the seismic data and reduced to a point cloud. Seismic data include a fair amount of redundancy in respect of geological information. Not every pixel of the seismic cube is necessary when conducting structural seismic interpretation. Instead, the interpreter aims to extract layers and features relevant to produce a mapping of the subsurface structures. The hypothesis in this thesis is that it is possible to filter the seismic cube to a set of points sufficient to conduct the desired geobody detection interpretation. The point cloud will ease the interpretation and data processing.

In conventional approaches, we treat the seismic cube as a 3D array. Most commonly, a Cartesian grid with coordinates of Inline (X-axis) number, Crossline (Y-axis) number and Time or Depth (Z-axis) sample. And every (x, y, z) position contains a property value. The property is the reflection amplitude or an attribute derived from the amplitude, such as coherence, envelope amplitude, instantaneous frequency, etc. In the point cloud approach, we treat the data as a 2D array (Equiv. to a table). Every line in this 2D array represents a point. The three first columns are the point coordinates (Inline value, Crossline value and Time or Depth value). The fourth column is the Amplitude value at the point. And any other attribute can be added as an additional column.

Point clouds are one of the most popular data structures used to represent 3D objects. It is used both as a processing format and as a visualization format. Many industries rely on point clouds for many purposes: to create 3D CAD (Computer-Aided Design) model in 3D printing [Kang et al., 2016, Mannoor et al., 2013], to represent and visualize landscape in Geographic Information System (GIS) [Shirowzhan and Sepasgozar, 2019, Szostak, 2020], for detection of surrounding environment for many self-driving car systems, or to render and animate images in entertainment with, for instance, the Kinect camera produced by Microsoft. The use of point clouds has intensified imaging and modelling outcrops in geology given the improvement and multiplication of 3D sensors with drones equipped with 3D scans.

Using a point cloud representation of seismic data opens up a whole field of active research on point cloud processing. The increasing use of Lidar data, particularly for self-driving cars, has intensified the research on the automated processing of 3D point clouds. As a result, significant progress has been made on the tasks of point cloud segmentation [Grilli et al., 2021, Poux and Billen, 2019, Poux et al., 2020, Qi et al., 2017] and shape recognition and retrieval [Lev, 2020].

3.2.2 Creation of a seismic data point cloud

The initial phase of the segmentation method aims to extract points relevant to the main structures inside the seismic cube. A seismic geobody is an interpreted 3D object which delineates geometric, structural and lithological patterns characteristic of a particular depositional feature. We interpret the geobody at the peak of the amplitude trace. The wavelength of the amplitude event will provide information on the thickness and volume of the geobody. A geobody will typically possess fairly similar amplitudes with continuous patterns. We selected the filtering processes creating the seismic point cloud based on these definitions of geobody.

In order to produce the seismic point cloud, the initial step consists in extracting every local extremum (minimum or maximum) of the amplitude in the trace direction (figure 3.3). The choice of extracting minima or maxima can be taken using prior information

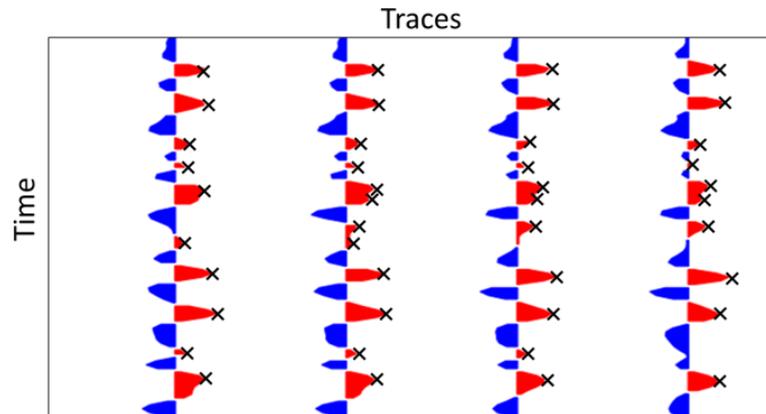


FIGURE 3.3: Extraction of local maxima (x) on seismic wiggle traces.

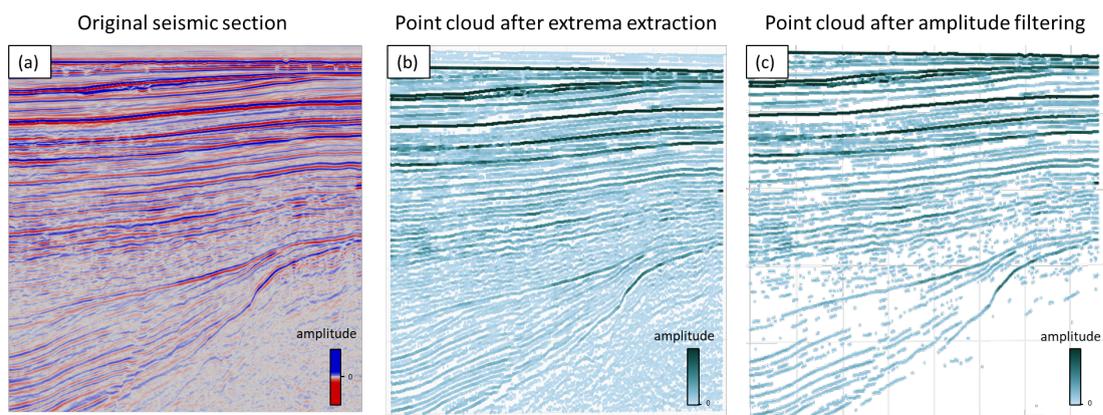


FIGURE 3.4: Creation of a Point Cloud Seismic of a seismic slice. (a) Input seismic slice. (b) Result after the extrema extraction to the slice, extracting only peaks. (c) Result after application of an amplitude cut-off to the remaining points of the extrema extraction.

known about the seismic response of the geobody - positive or negative polarity event. If no prior knowledge of peak or trough events is available, we recommend extracting both minima and maxima. The seismic amplitude measures the relative contrast of the impedance of the rock. Conventional seismic visualisation displays the full seismic waveform record, but one single extremum point suffices to describe the seismic reflectors [Borgos et al., 2003, 2005]. This trace peak or trough extraction could be improved by fitting a polynomial curve to the trace in the seismic grid before performing the extrema extraction. Horizon interpretations are often placed on the extrema events in the seismic cube. Figure 3.4 displays the results of the point cloud extracted from an example seismic slice. The extrema extraction preserves every feature displayed in the slice while reducing the number of points and producing space between events.

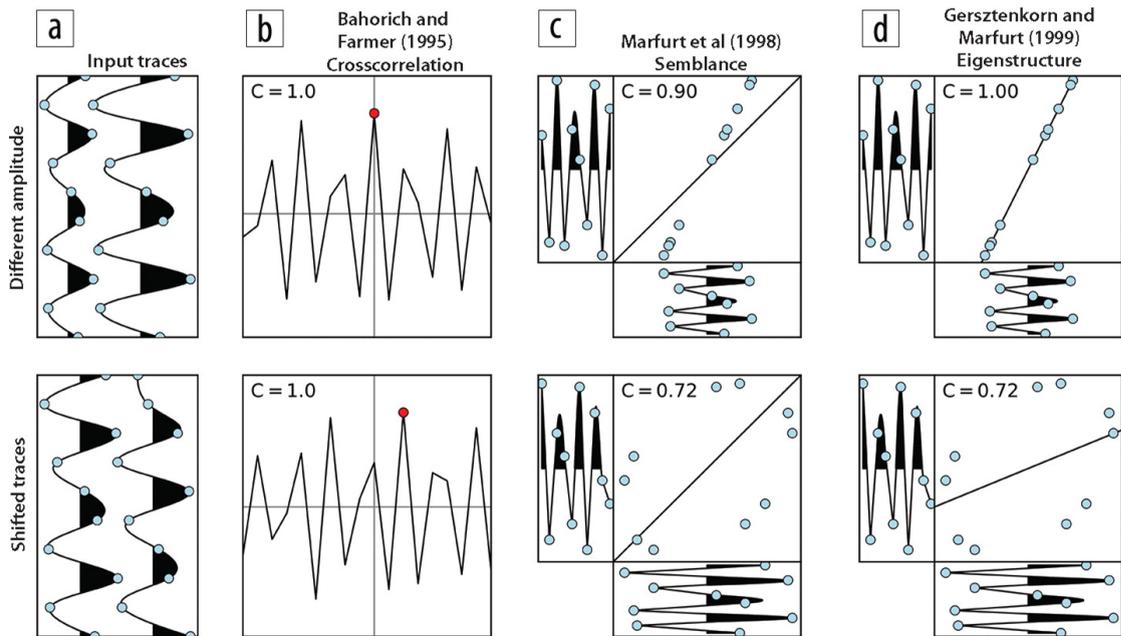


FIGURE 3.5: Visual explanation by Kington [2015] of the differences among the coherence implementations discussed using only two traces. Letter C indicates the calculated coherence value.

The second step is to filter the points based on their coherence value.

Coherence is one of the most commonly used seismic attributes. Coherence is part of a set of post-stack calculations that measure seismic traces' local similarity or dissimilarity. Low similarity zones appear where there are discontinuities in a seismic volume, such as faults or mass transport deposits [Wu, 2017]. As a result, we can use discontinuity attributes to indicate faults in a seismic volume or subtle stratigraphic structures. Many terms are used to refer to attributes measuring dissimilarity: coherence, semblance, similarity or discontinuity. Bahorich and Farmer [1995] first developed coherence. The implementation is based on the maximum cross-correlation value of three seismic traces. Marfurt et al. [1999] extended Bahorich and Farmer's cross-correlation technique to any number of input traces, thus reducing the sensitivity to noise. Gersztenkorn and Marfurt [1999] improved the implementation by computing the ratio of the largest eigenvalue of the covariance matrix to the sum of eigenvalues, making it more robust to lateral amplitude differences. The details and comparisons of these different implementations are available in Kington [2015] and illustrated in figure 3.5.

When extracting points for the seismic point cloud representation, the idea is to take

advantage of the high amplitude consistency of the geobodies when possible. By keeping only the points with high coherence values, we preserve the geobodies and filter the discontinuous points. Thus, we create the distance between the seismic features. Especially at fault zone locations, the fault offset is likely to have put side-by-side unrelated seismic events. By filtering on coherence, we create spaces between these non-consistent point cloud events.

Additionally, we apply a cut-off value on low amplitude values to preserve structures with the highest impedance contrast. The filtering of low-amplitude values makes it possible to eliminate the noisy points that remain after the extrema extraction and do not translate directly into geological objects. For example, in figure 3.4, we apply a simple amplitude cut-off to get the major stratigraphic events and suppress non-desired features such as basement response and secondary structures.

The choice of the cut-off values results from a trade-off between filtering out noise and preserving as many structures as possible. An intuitive approach to finding a good cut-off value is to select and visualize the results of the cut-off on a slice of the seismic (figure 3.6). Generally, a value corresponding to the first third of the points produces a satisfactory result.

Another way to think about these cut-off values is to see them as uncertain parameters that can be inferred in a Bayesian Uncertainty Quantification framework. Indeed, each parameter set leads to a different output, ultimately leading to a different seismic segmentation and geobody detection. So, each of these outputs can be described as an additional possible interpretation scenario. This idea is further explored in the discussion of chapter 4.

3.2.3 Application of the point cloud extraction to the *validation dataset*

We applied the extraction of the point cloud seismic to the *validation dataset* described in *Chapter 1 Research methodology*. The *validation dataset* consists of 1171 inlines, 1714 crosslines and 708 time-samples, so 1,421,022,552 voxels (10^9). The first step was the

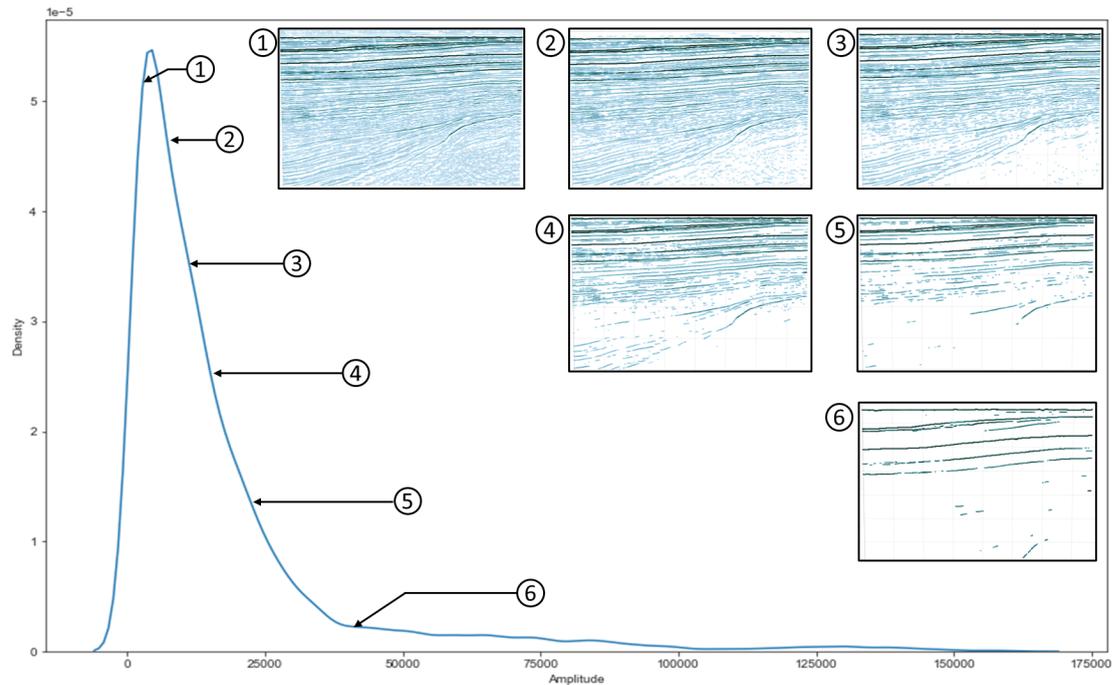


FIGURE 3.6: Application of different amplitude cut-off values to a seismic slice from 1 to 6. The selected cut-off values are indicated on the kernel density estimation of amplitude of the points of the seismic slice.

extrema extraction. We extracted only maxima here because the focus was on turbidite fans, and fans are peak events in that seismic. The extrema extraction took 8min38s to execute and reduced the seismic data to 169,779,498 points (10^8). Experiments were run on a laptop with an Intel Core i9-10885h CPU processor and 32GB of RAM, these specifications are used for all results produced for this thesis. The second step was the filtering of coherence. Computing coherence and filtering took 3min21s to run and reduced the point cloud to 101,581,924 points (10^8). The third step was the filtering of amplitude. Filtering amplitude took 2s to execute and reduced the point cloud to 43,201,714 points (10^7). So that makes a total time of 12min01s for the seismic point cloud extraction and a reduction factor of 97% from voxels to points.

The first evaluation of the point cloud is visual. A few quick observations are possible by plotting the entire point cloud in a 3D interactive view (figure 3.7). Firstly, the basement response has been filtered out on the entire cube. Secondly, we can quickly identify two main types of structure stacks: long extended horizons that cover the whole area on the shallower part of the point cloud and, in the deeper part, some more constrained deposits filling the paleolake, among which are some turbidite fans coming from the slope. At

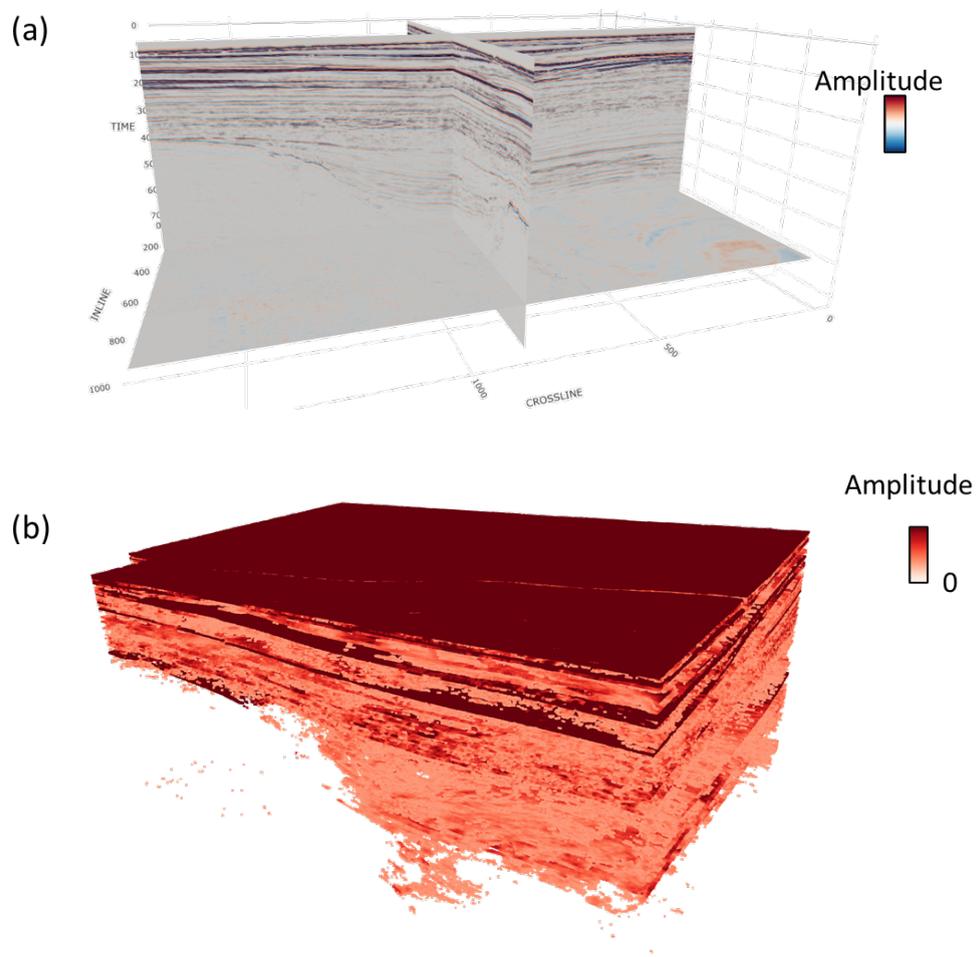


FIGURE 3.7: (a) section view of the 3D seismic used as *validation dataset*; (b) point cloud extracted after extrema extraction, coherence filtering and amplitude filtering.

least the shallower horizons are well preserved. We can observe their full extent, even the canyon on the top horizon (figure 3.7b). In the fillings of the paleolake, hundreds of seismic features are present. By focusing on the area of the point cloud where two turbidite fans are, we can observe that they are also well preserved (figure 3.8). The channel feeder and the lobes structure of the turbidites are easily detectable. All these observations validate the extraction of the point cloud from the seismic data.

3.2.4 Advantages and Disadvantages of a seismic data point cloud

The point cloud format brings three main advantages:

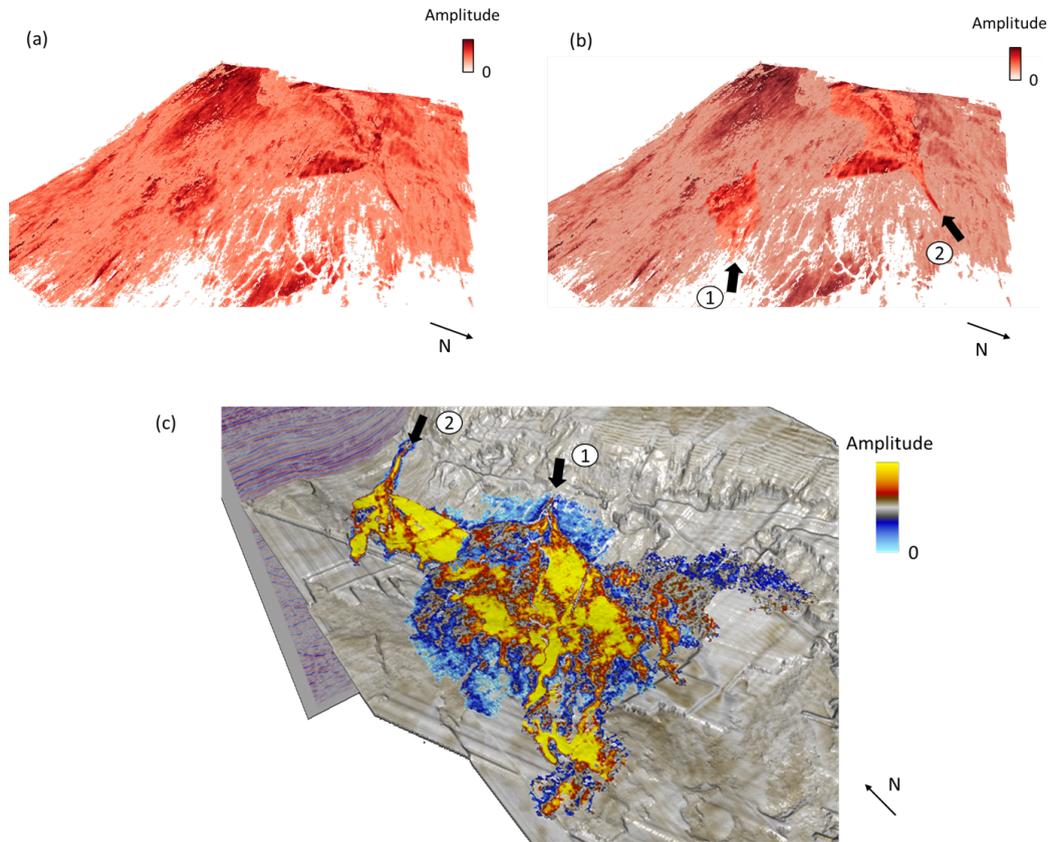


FIGURE 3.8: (a) Visualisation of the seismic point cloud, focused on the area where the turbidite fans are. (b) Same visualisation with the two turbidite fans highlighted. The black arrows point to the fan input feeders. (c) The two turbidite fans extracted using Petrel seed-based horizon auto-tracking. The color scale has been selected to better display the amplitude texture.

- (i) A sparse data allows 3D visualisation without being constrained to 2D seismic sections. We take a small 3D example as an example (figure 3.9) to illustrate the advantages of visualising the seismic as a point cloud. The sparse representation of the data allows visualising and understanding of the structures in a glimpse. We observe the extent and delineation of the horizons and the displacement along fault planes.
- (ii) It highlights the distances between objects in the seismic data. Furthermore, the possibility of working on distances between points opens new ways of processing the data. For example, in a point cloud, it is possible to work on the spatial continuity of structures rather than being limited to image and texture analysis.
- (iii) It reduces the amount of data by a factor from 10 up to 100. Thus, it improves the computation efficiency and memory requirements for data processing. Moreover,

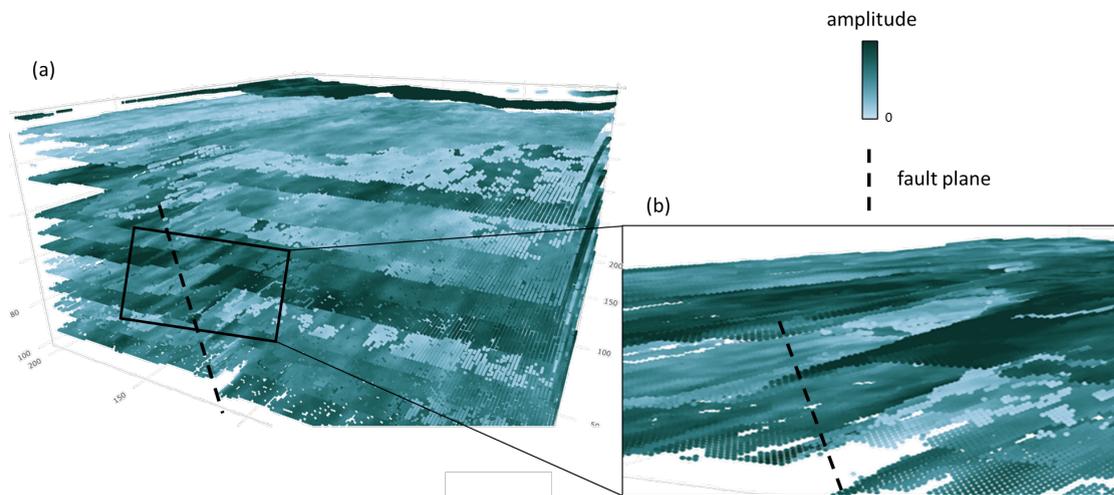


FIGURE 3.9: (a) Point Cloud Seismic result of a small 3D seismic cube. (b) Zoom in to a particular zone of the point cloud highlighting the displacement of horizons along a fault plane.

sampling is straightforward and preserves structures' distance and relative density.

On the downside of the point cloud format:

- (i) Filtering data represents several risks of loss of information. The first risk of loss is to filter out a structure or a geobody that would be of interest. Applying only low thresholds on amplitude and coherence should reduce that risk. But there is no guarantee that one will not filter a geobody with a weak seismic response or in a noisy part of the seismic. The second loss of information is the loss in terms of fine details that the full trace wavelet has. One way to mitigate the loss is to tie local amplitude texture and wavelet characteristics to the extremum point extracted. We compute the wavelength and the wave area and store these attributes as additional properties to every point of the point cloud.
- (ii) The other downside or constraint is the complexity that goes with the processing of point cloud data. As a non-regular structure in the form of a list of points, point cloud analysis is complex. Therefore, applying convolution with kernel operations to point clouds is not straightforward compared to regular grids. However, recently, algorithms including graph kernels [Bach, 2008] have been developed and allow, in particular, kernel approaches for classification and segmentation tasks [Lin et al., 2020]. Additionally, because the set of points is unordered and with a

variable number of points, any method applied to the point cloud must be robust to data permutation and size [Lev, 2020]. This processing complexity is the main reason why there have been few works on point clouds until recently.

3.2.5 Visualisation and plotting of 3D Seismic as a Point Cloud

In order to take advantage of the sparsity of point cloud data, one might want to visualise the full sparse data cube at once. For a regular size of a 3D seismic, the remaining number of points in a seismic point cloud will typically be between 10 million to 100 million points. However, plotting and rendering 100 million points in an interactive viewer is extremely demanding and will likely exceed the computer's available memory capacities. One easy solution is to sample the dataset and plot only a subset of points. The sampling can be performed randomly or voxel-based to preserve the homogeneity of the distribution. But sampling is not ideal because it reduces the level of details that get visualised. A better solution is to use out-of-core real-time rendering techniques [Richter and Döllner, 2010, 2014]. Out-of-core rendering systems use multiple Level-of-Detail (LoD). It uses a tree data structure that hierarchically subdivides the spatial area [Wimmer and Scheiblauer, 2006]. Each LoD is represented by a node of the tree (figure 3.10). Figure 3.11 illustrates four different LoD applied to an example point cloud. In real-time rendering, the user interaction and the view position trigger a selection of different LoD nodes in each visualisation frame. The groups of points far from the user get approximated as single points, and the points outside the view frustum are removed. It reduces the number of points that needs rendering. Hence it allows overcoming the limit of memory, making visualisation and interaction of a massive 3D point cloud feasible.

Recent advances in Virtual Reality (VR) are leading to the development of applications with more immersive visualization. The immersive VR experience allows interacting with the data differently. However, moving in a virtual environment implies a sparsity of the data so the user can see in the distance. So preferable formats are meshes or point clouds. Hence, treating seismic data as a point cloud opens perspectives on easily integrating seismic in VR applications without much preprocessing or interpretation.

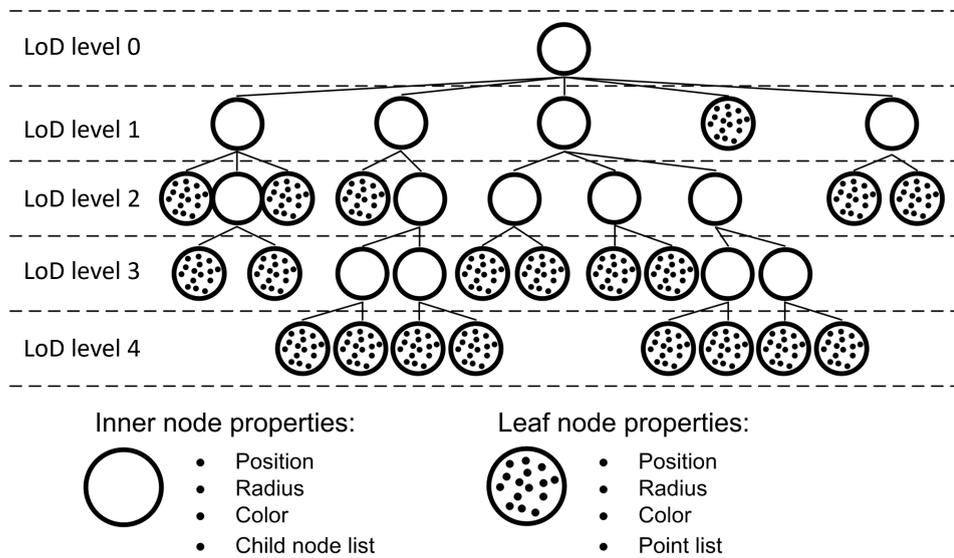


FIGURE 3.10: Octree structure which subdivides spatially the point cloud into different Level-of-Detail (LoD). Figure from Richter and Döllner [2014].

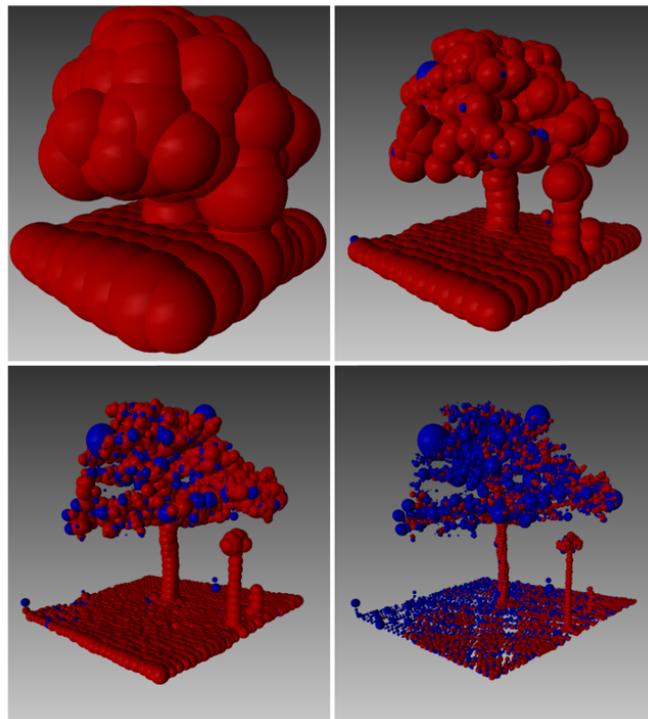


FIGURE 3.11: Illustration of the tree structure with bounding spheres for 4 different LoD. Each sphere represents a tree node and encloses all child nodes of the presented node. Inner nodes are illustrated in red and leaf nodes in blue color. Figure from Richter and Döllner [2014].

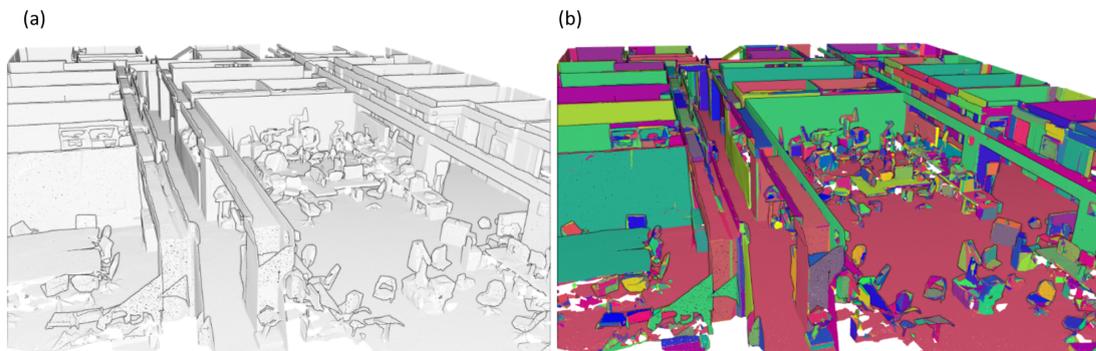


FIGURE 3.12: Illustration of the planar segments produced for building interior classification. (a) Point cloud; (b) Segmentation. Figure from Bassier et al. [2020].

And geology, with its high diversity of data structures (seismic, wells, maps...), could benefit from the development of the immersive visuals of VR.

3.3 Segmentation of a seismic data point cloud

3.3.1 Approaches to point cloud segmentation in other industries

Since point cloud is a popular data format, many studies have been conducted on point cloud segmentation for other types of applications and industries. Amongst the first historical approaches to point cloud segmentation, a frequently used technique is to over-segment the point cloud into small regions before applying an expensive algorithm to create more significant regions. The most straightforward strategy is to use voxels as an over-segmentation structure and then perform voxel-clustering [Son and Kim, 2017, Weber et al., 2010]. The emergence of deep learning approaches improved the performance of point cloud segmentation [Maturana and Scherer, 2015, Qi et al., 2017, Zhang et al., 2019] and, more particularly, semantic segmentation. However, such algorithms depend on annotated datasets for training, requiring operators with high-level expertise to annotate large quantities of data [Grilli et al., 2021]. On the other hand, unsupervised clustering approaches do not require such training data and have also been successful, especially for interior classification [Grilli et al., 2021, Poux and Billen, 2019, Poux et al., 2020] (figure 3.12).

However, such unsupervised solutions are designed around the particular properties of building interiors and cannot be directly applied as such to a seismic point cloud. For instance, a commonly used algorithm is the RANdom SAmple Consensus (RANSAC) [Schnabel et al., 2007]; RANSAC segments the point cloud by planar structures, which fit well for interior segmentation but not geological objects.

3.3.2 Density-based Segmentation

In the next step of our seismic segmentation workflow, we segment the seismic data point cloud into distinct objects. We propose to perform the segmentation by clustering the point cloud based on the spatial density of the points. The intuition behind this approach is to take advantage of the distance between the seismic events created by the point cloud representation to segment them spatially. We use Density-Based Spatial Clustering of Application with Noise (DBSCAN) to perform this point cloud segmentation.

DBSCAN is an unsupervised clustering algorithm relying on a density-based notion of clusters [Ester et al., 1996, Schubert et al., 2017]. Every point is classified into three categories: core point, border point and noise, based on the density of their neighbourhoods (figure 3.13a).

For each point, DBSCAN evaluates the density of its neighbourhood at an Epsilon distance (Eps). The Eps -Neighbourhood is the ensemble of points included at a distance inferior to Eps in the three spatial directions. The Eps -Neighbourhood of a point p ($NEps(p)$) is defined by the equation 3.1:

$$NEps(p) = \{q \in D \mid dist(p, q) \leq Eps\} \quad (3.1)$$

If the $NEps(p)$ contains at least a minimum number of points ($MinPts$), the point p is classified as a core point. A point in the Eps -Neighbourhood of a core point but not dense enough to be a core point is classified as a border point. A point neither dense enough nor in the Eps -Neighbourhood of a core point is classified as noise. Then, the definition of clusters in the DBSCAN algorithm relies on the notions of density reachability and

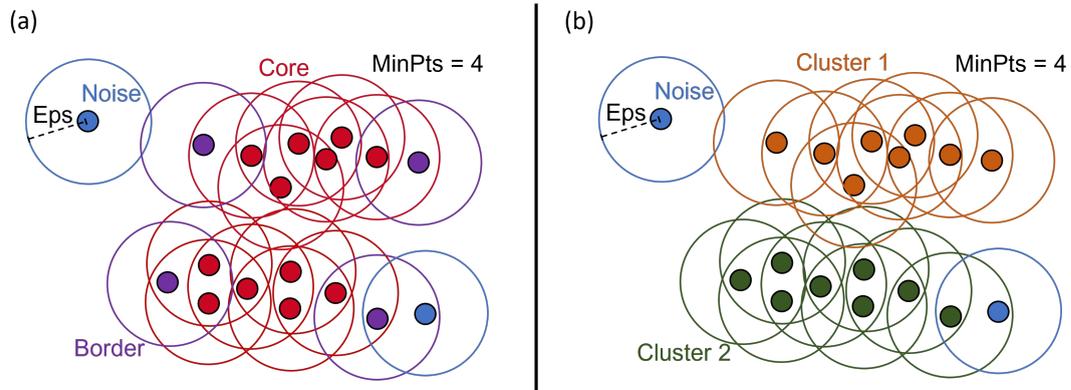


FIGURE 3.13: This diagram illustrates clustering of points with DBSCAN in a 2-dimensional space. The circles around points materialize the limits of the Eps-Neighbourhood of the points. Figure (a) presents the classification into types of points. The red points are core points; they have an Epsilon-neighbourhood of at least 4 points (*MinPts*). The purple points are border point; they are reachable by core points but do not have an epsilon neighbourhood of at least 4 points. The blue points are noise; they are not reachable by core points. Figure (b) presents the resultant clustering performed from the point classification. A cluster regroups the ensemble of density-connected core and border points. Noise points remain as outliers from the clusters.

connectivity. Two points are directly density-reachable if they are included in the *Eps*-Neighbourhood of one another and if at least one of them is a core point. Two points p_1 and p_n are density-reachable if there is a chain of points $p_1 \dots p_n$ such that p_{i+1} is directly-reachable from p_i . Two points p and q are density-connected if there is a point o such that p and q are density-reachable from o . We define a cluster in DBSCAN as the ensemble of all the points density-connected, which corresponds to the ensemble of connected core points and border points. Finally, all the points not density-reachable by a core point are set as noise and remain outside the clusters (figure 3.13b).

DBSCAN offers several advantages:

- (i) its suitability to create spatial clusters,
- (ii) its capacity to handle noise,
- (iii) its control of the clustering with only a few parameters
- (iv) its simplicity and efficiency of its implementation – Open3D implementation of DBSCAN is used in [Zhou et al., 2018] –.

Point distributions of seismic events in the seismic point cloud are mainly dense points with space at the extremity between the events. DBSCAN algorithm allows the creation of clusters of various shapes of points based on their density-connectivity. That density-connectivity approach directly fits the task of spatial segmentation of a seismic data point cloud. Applying the clustering to the 3-dimensional space of the seismic data point cloud creates spatially consistent clusters from the spatial point distribution. Furthermore, seismic data are noisy data, so the ability of DBSCAN to leave outliers unclustered adds robustness to the segmentation. Compared to other clustering algorithms, DBSCAN does not requires inputting any prior information on the expected shapes of the clusters or the number of clusters that are information unknown at that stage. Instead, we control the clustering by only two parameters: an epsilon distance on which we evaluate the density of points and a minimum number of points which sets as a threshold on the density to define connected points. With the meaning of these two parameters in mind, one can quickly test and select the value of these parameters to get the desired segmentation. Finally, DBSCAN is suitable for implementations using kd-trees [Bentley, 1975] and balltrees [Omohundro, 1989] to determine the neighbourhood of points, which avoids calculating the total distance matrix and makes it very computing efficient.

To better perform the segmentation of the seismic data point cloud, we propose to add a third parameter, a vertical exaggeration factor ($Zexg$). Seismic horizons and geobodies are often laterally elongated structures rather than vertically. Therefore, multiplying the point cloud's vertical value by $Zexg$ accentuates the vertical distance between the points (figure 3.14). As a result, when segmenting the vertically exaggerated point cloud with the same Eps in the three spatial dimensions, the connection between the points is laterally stronger than vertically. Thus, the segmentation better accommodates the lateral extent of the seismic features (figure 3.15).

The choice of DBSCAN hyperparameters – Eps , $MinPts$ and $Zexg$ – mainly depends on the density of the point cloud. The parameters must be adjusted to each other. As a general principle, the higher Eps , the wider the structures will be. On the other hand, $Zexg$ has to be strictly inferior to Eps to allow the clustering of points with vertical offsets. These distances have to fit the interpretation target and depend on how extensive and connected the geobodies are. After choosing the values for Eps and $Zexg$, the

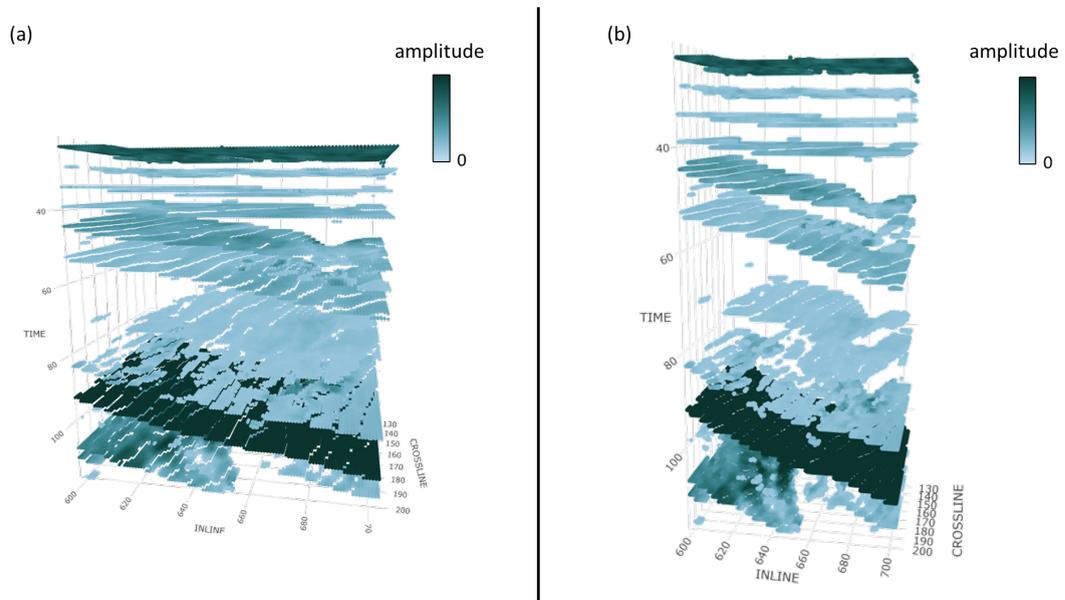


FIGURE 3.14: Vertical exaggeration of a seismic data point cloud for $Zexg=1$ (a) and $Zexg=2$ (b).

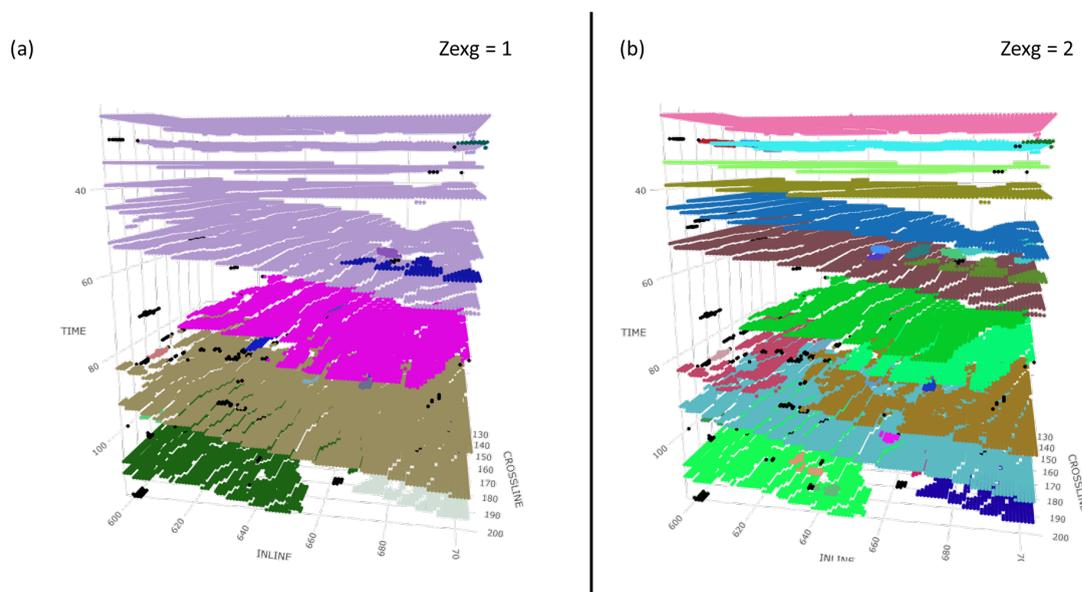


FIGURE 3.15: Illustration of the impact of $Zexg$ on the segmentation with DBSCAN for $Zexg=1$ (a) and $Zexg=2$ (b). Other DBSCAN parameters (Eps & $MinPts$) are identical for both clusterings. Every cluster is associated with a random color.

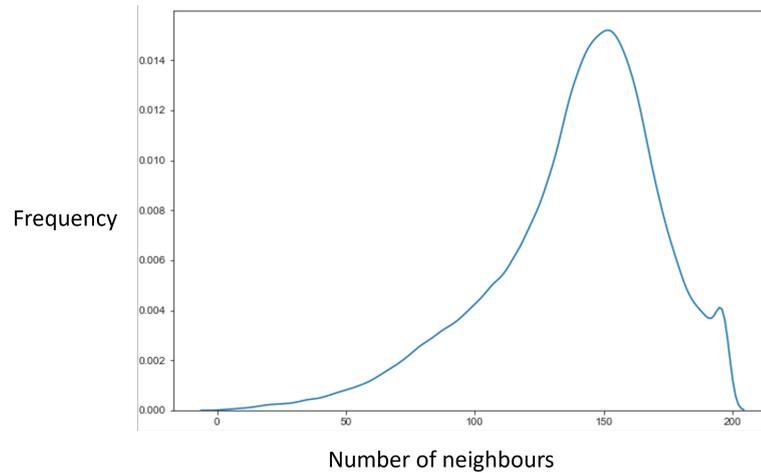


FIGURE 3.16: Kernel density distribution of the number of neighbours in the epsilon-neighbourhood of the points for $Eps=8$. It indicates a value between 50 and 100 to be considered as $MinPts$ given the selected Eps .

question is: how many points lie within each point's epsilon-neighbourhood? A good practice is plotting the number of neighbours within each point's epsilon-neighbourhood for the selected Eps (figure 3.16). Based on Figure 3.16, most of the points (>75%) have more than 100 neighbours. A small fraction has less than 50 neighbours; they probably are noise points. We must adjust the value on how noisy the data is. Here, we could consider a value between 50 to 100 for $MinPts$.

Another interesting clustering algorithm for the segmentation of the seismic point cloud is Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN). HDBSCAN was developed by Campello et al. [2013] and was inspired by DBSCAN. In McInnes et al. [2017]'s implementation, HDBSCAN operates similarly to DBSCAN to transform the space according to the density of the points. It computes the mutual connectivity between the data points by progressively increasing the epsilon value. All these mutually connected data points form a graph called minimum-spanning-tree for mutual reachability distance. We represent this tree as a dendrogram (figure 3.17). The dendrogram is obtained by sorting the edges of the tree by distance and creating a new merged cluster for each edge. We can also describe this dendrogram as DBSCAN clusters varying epsilon values integrated into a hierarchical tree of the different cluster levels. This hierarchical tree is then condensed to a simpler tree of persistent clusters. This more condensed tree representation is based on the notion of minimum

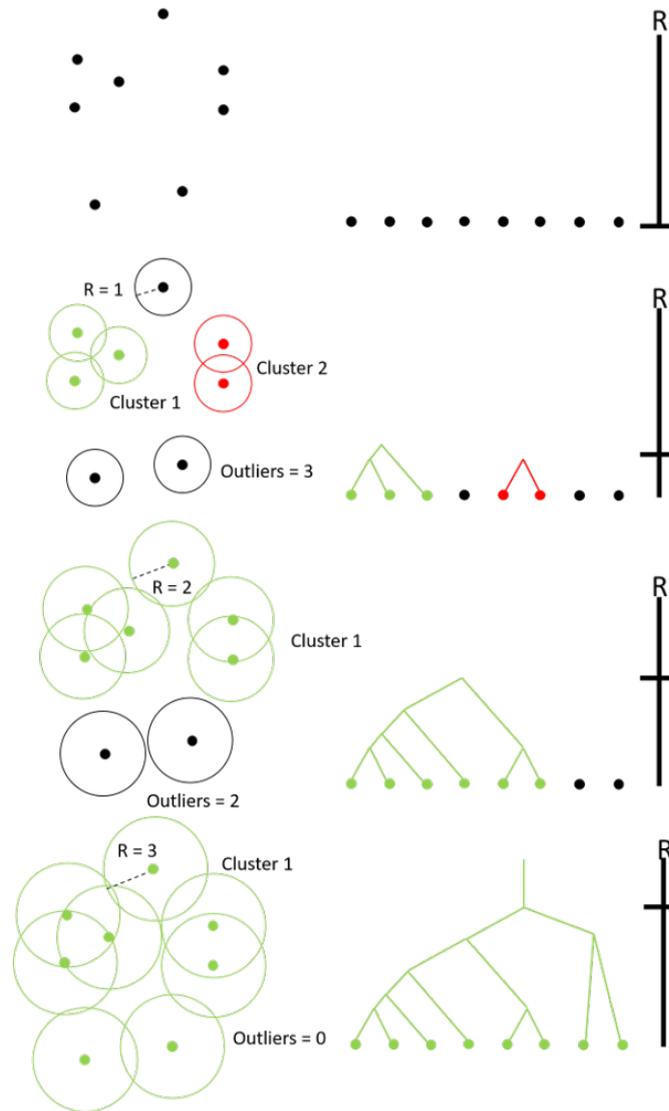


FIGURE 3.17: Gradually increasing radius r causing the generation and merging of new clusters. This can also be visualised with the help of a dendrogram. From Steffens [2022].

cluster size (min_cl_size). min_cl_size is a hyper-parameter of HDBSCAN. Starting from the top of the dendrogram, at each split, a conditional statement is performed regarding the size of the new clusters produced. If one of the new clusters produced by a split has fewer points than the min_cl_size , these points are considered as falling out of the cluster. But, on the other hand, if the two clusters have at least the min_cl_size , then the split persists in the tree. In this condensed tree representation, the width of the line represents the number of points in the cluster. The condensed tree is expressed using lambda, the inverse of the epsilon distance: $\lambda = \frac{1}{distance}$. From this condensed tree representation, we extract the most significant clusters. For each cluster in the condensed

tree, we compute the stability as:

$$Stability_{cluster} = \sum_{p \in cluster} (\lambda_p - \lambda_{birth}) \quad (3.2)$$

Where λ_p is the lambda value at which that point fell out of the cluster, and λ_{birth} is the lambda value at the split at which the cluster starts. We can also view that representation as a measure of cluster stability over the varying epsilon distances in the DBSCAN framework [McInnes et al., 2017]. The points that fell out of the tree at higher hierarchical positions than any selected stable cluster in their parent splits are considered noise. We illustrate how HDBSCAN works on an example dataset in figure 3.18.

Given a proper hyper-parameter selection, HDBSCAN ($min_cl_size = 200$) and DBSCAN ($eps = 2.0$, $MinPts = 8$) produce a segmentation of the seismic point cloud of similar quality. Figure 3.19 illustrates the segmentation performed by both algorithms on a small cube. It shows the best realisation of both solutions obtained with a minimal effort to tune the hyper-parameters - 5 minutes of manual tuning of the parameters. DBSCAN leads to 2,186 clusters and HDBSCAN to 273 clusters only. However, the difference is due to small clusters preserved in DBSCAN result. Indeed, only 237 clusters have more than 200 points. Both algorithms are based on the same approach of density-connectivity between the points of the same clusters. And both algorithms lead to similar results: all the main clusters of both segmentation in Figure 3.19 are identical, whether it is a wide horizon or a more constrained body.

HDBSCAN offers several positive arguments for the segmentation of the seismic cube. Compared to DBSCAN, HDBSCAN relies on a more intuitive parameter: the minimum cluster size is more straightforward to apprehend than the estimation of the point density coming from the two parameters of DBSCAN - the epsilon distance and the minimum number of points. HDBSCAN also offers the possibility of having varying density clusters. Moreover, HDBSCAN adds a hierarchical component to the seismic segmentation, which applies well to the notion of segmentation of geological objects (figure 3.20). In the example of a turbidite fan, we could segment a fan into lobes and a channel feeder, or we could segment it at the turbidite fan level, or we could segment it together with

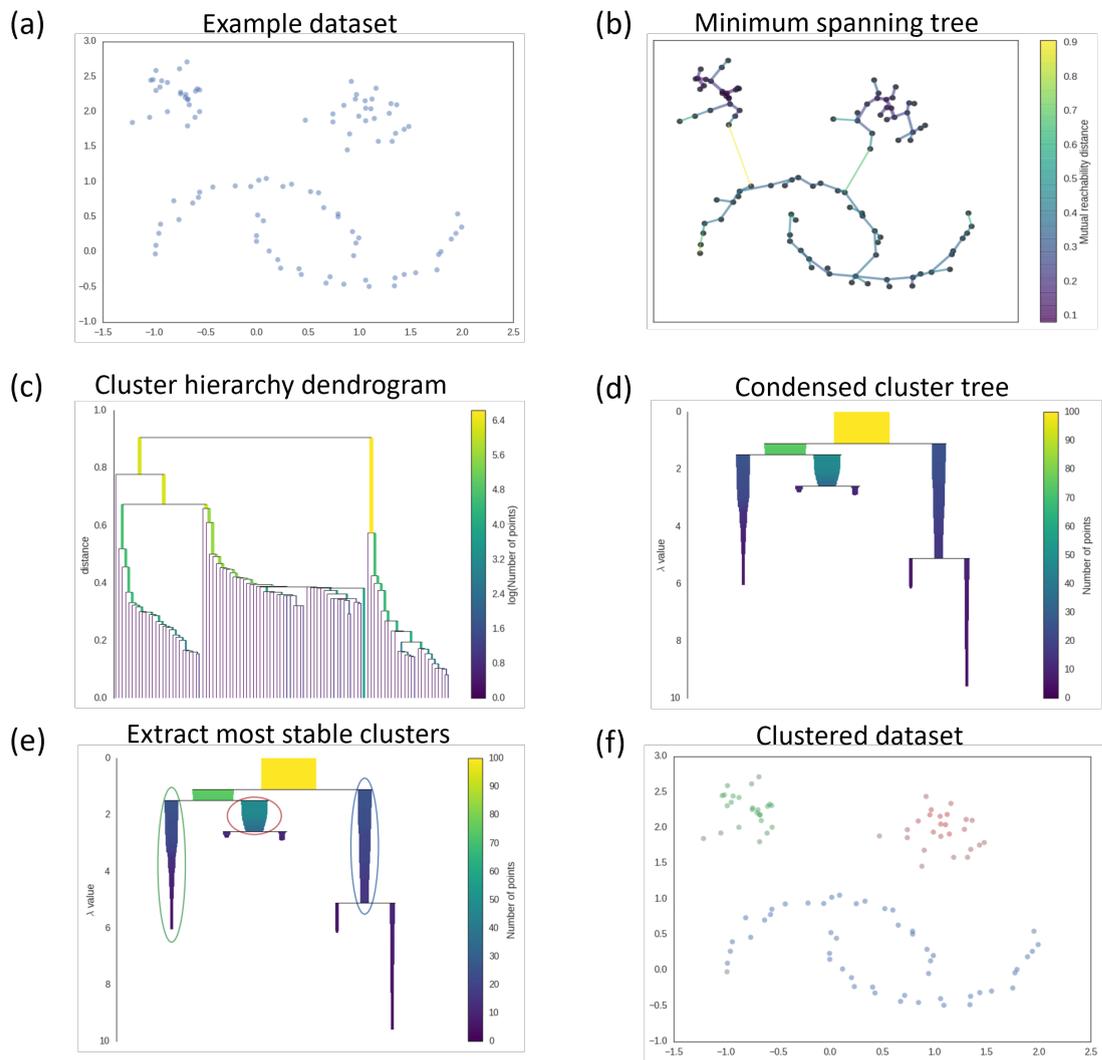


FIGURE 3.18: Illustration of the different operations performed by HDBSCAN on an example dataset to extract clusters. Adapted from McInnes et al. [2017] and HDBSCAN documentation [McInnes, 2022].

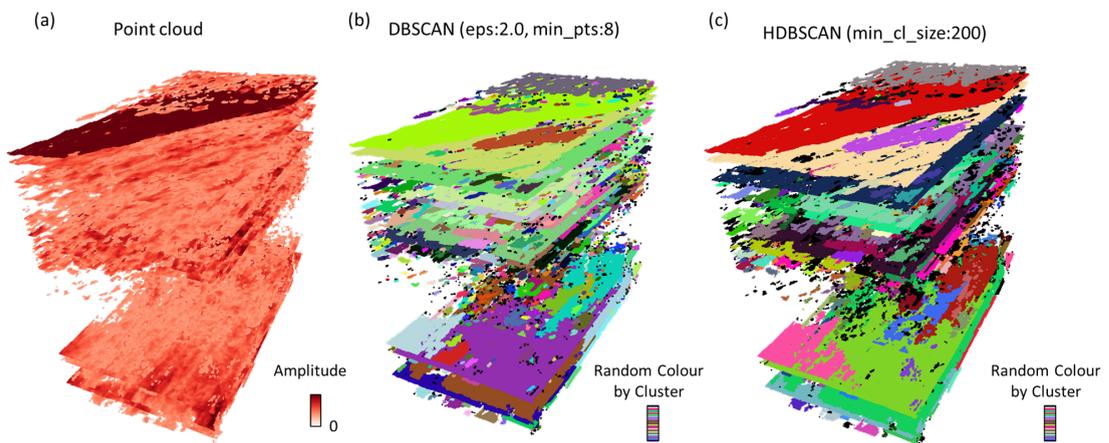


FIGURE 3.19: Comparison of DBSCAN and HDBSCAN results on the segmentation of a small 3D cube.

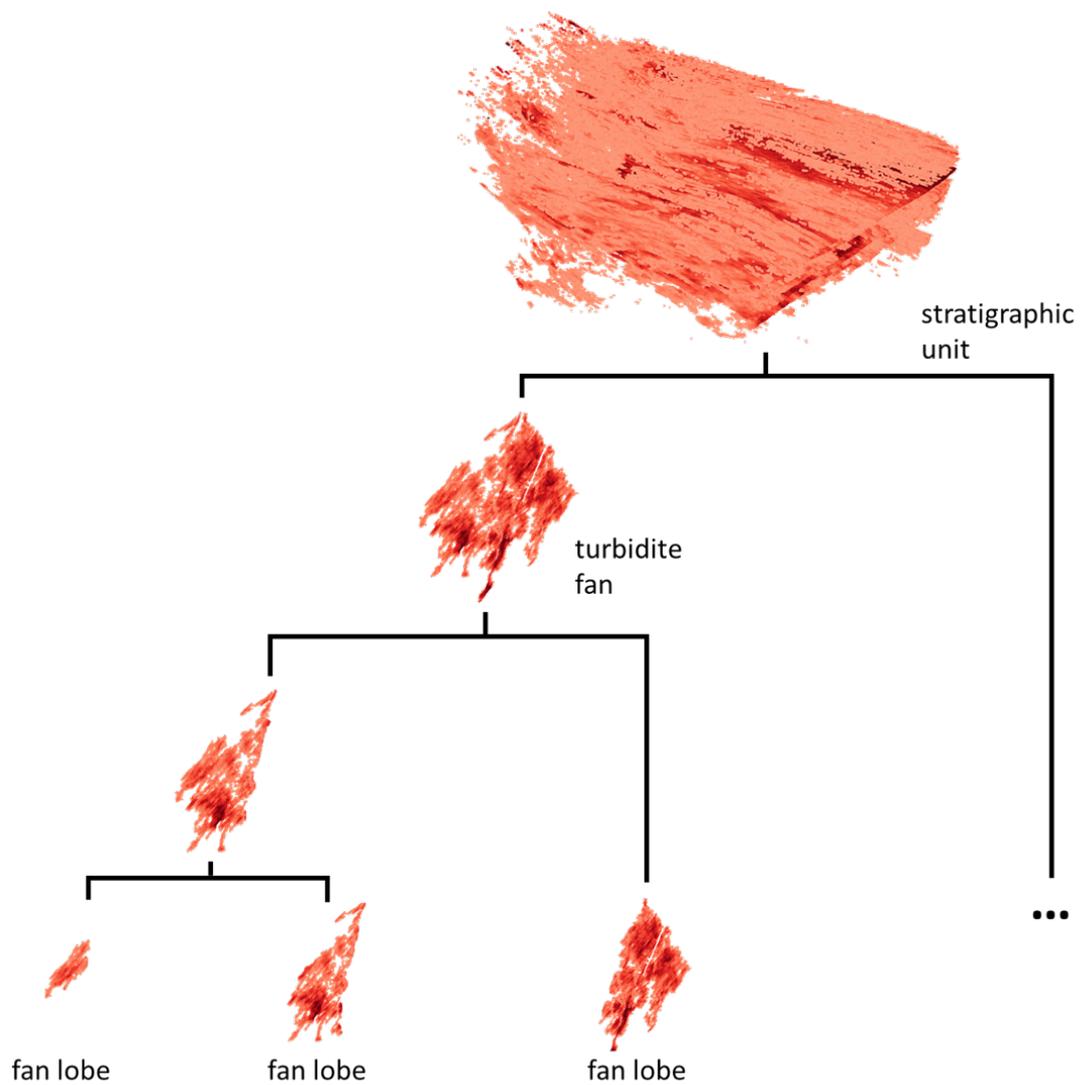


FIGURE 3.20: Illustration of the notion of hierarchy in the geologic structures.

its stratigraphic unit. Selecting the cluster level on a hierarchical tree would provide adaptability to the segmentation. Thus, it would allow adapting the segmentation to the interpretation task.

However, several arguments favour the choice of DBSCAN rather than HDBSCAN. Firstly, working on a specific connectivity distance is better when working on the segmentation of a seismic cube. If having varying density clusters is beneficial for other applications, this is not the case for the segmentation of seismic point clouds. We obtain the seismic point clouds from data evenly sampled in the 3D space. Therefore, it is better to apply the same minimal distance delimiting connectivity of the clusters

of points. Thus setting a specific epsilon-connectivity for the entire cube is better than basing it on the cluster stability over varying epsilon distances. The cluster stability could merge subtle features into wider, more stable horizons. HDBSCAN, in its current implementation, does not allow the interpreter to control the clustering level. Secondly, DBSCAN scales better to large datasets. McInnes and Healy [2017]’s implementation of HDBSCAN offers a significant performance improvement, bringing to approximately a $O(N \log N)$ time complexity. But still, HDBSCAN does not scale well to a large dataset such as the 3D point clouds we intend to process for seismic application. Ester et al. [1996]’s implementation of DBSCAN is more adapted to large datasets with several millions of points. For instance, the segmentation of 752,841 points used for Figure 3.19 took 1.6s for DBSCAN -Zhou et al. [2018]’s implementation- and 54s for HDBSCAN -McInnes and Healy [2017]’s implementation-. And this execution time difference will only increase with bigger seismic cubes. Time performance being a requirement for the practicality of the proposed method, DBSCAN is advised and will be used for segmentation throughout the rest of this thesis.

3.3.3 Application of the segmentation to the *validation dataset* point cloud

We now apply the segmentation to the point cloud obtained from the *validation dataset* and described in section 3.2.3. After some initial tests to determine acceptable parameter values, we choose an epsilon value of 2.5 and a minimum number of points of 16 to perform the segmentation with DBSCAN. We do not add a vertical exaggeration to the segmentation. The segmentation runs in 4 minutes for a point cloud of 43,201,714 points. The segmentation result is displayed in Figure 3.21. The segmentation contains 60,388 clusters. The biggest cluster has 4,269,134 points, and only 3% of the clusters have more than 1,000 points (1,605 clusters). The number of noise points is 1,614,360. So, it represents a noise ratio of 3.7%, which is an acceptable proportion.

Figure 3.22 focuses on the impact of the segmentation on the two turbidites of interest in the *validation dataset*. As desired, we observe that the two fans are segmented into two clusters. The segmentation is good: the two fans are neither associated in the same

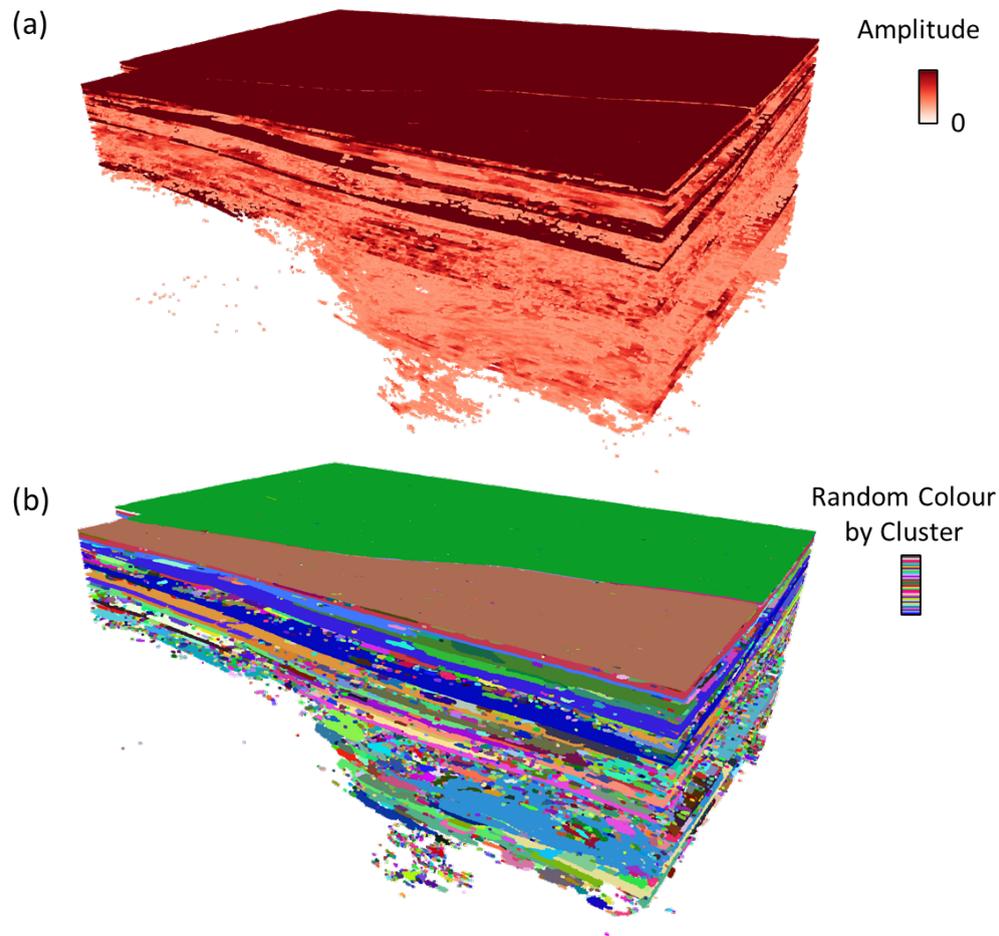


FIGURE 3.21: Results of the segmentation with DBSCAN ($Eps:2.5$, $MinPts:16$) on the point cloud obtained from the *validation dataset*.

cluster as other surrounding structures (under-segmentation) nor split into several clusters (over-segmentation). Even the more isolated points, such as the ones that constitute the feeder channels, are nicely preserved. However, by comparing the turbidite fan extracted with the point cloud segmentation to the turbidite fan extracted with Petrel seed-based auto-tracking (figure 3.23), we notice that low amplitude elements at the edges of the geobody are not preserved. It is probably due to the amplitude cut-off. Also, some less connected parts of the fan are missing at the South-East and South-West. Fine-tuning the amplitude filtering and segmentation parameters could improve the results, preserving lower amplitude and merging lower connected bodies during the segmentation. But the eventuality of missing a part of a geobody in the segmentation is a potential limitation to the point cloud segmentation method.

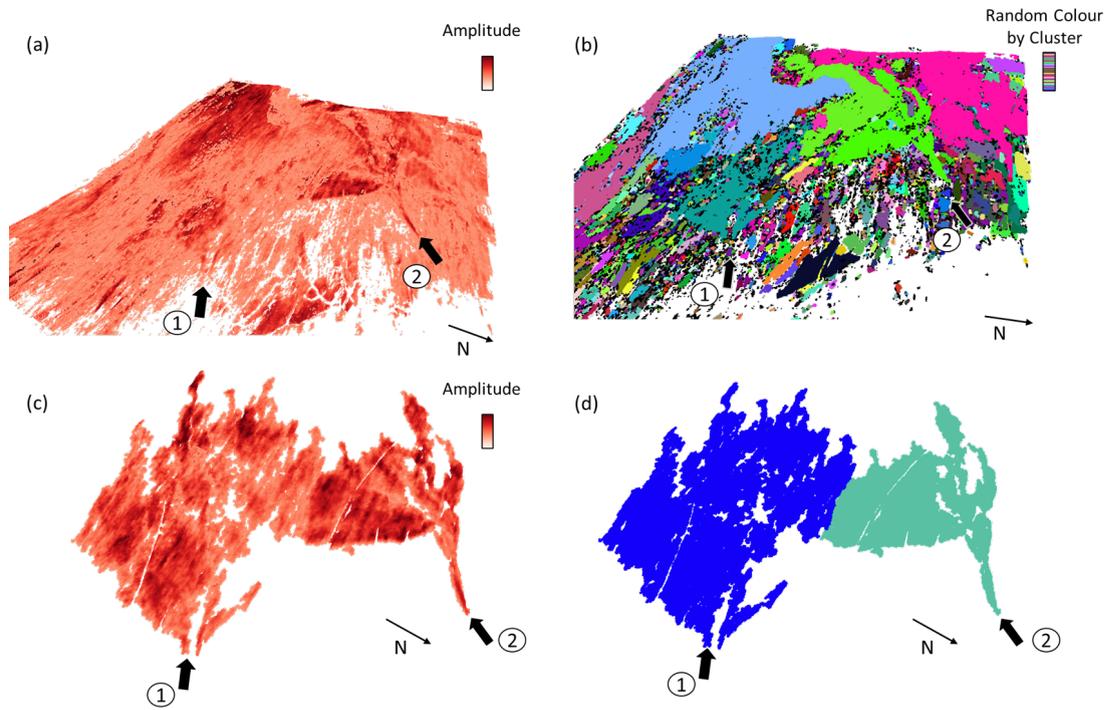


FIGURE 3.22: Results of the segmentation focused on the area where the turbidite fans are. (a) and (b) show the two turbidites among the surrounding structures recorded in the seismic data. (c) and (d) show the two turbidites extracted.

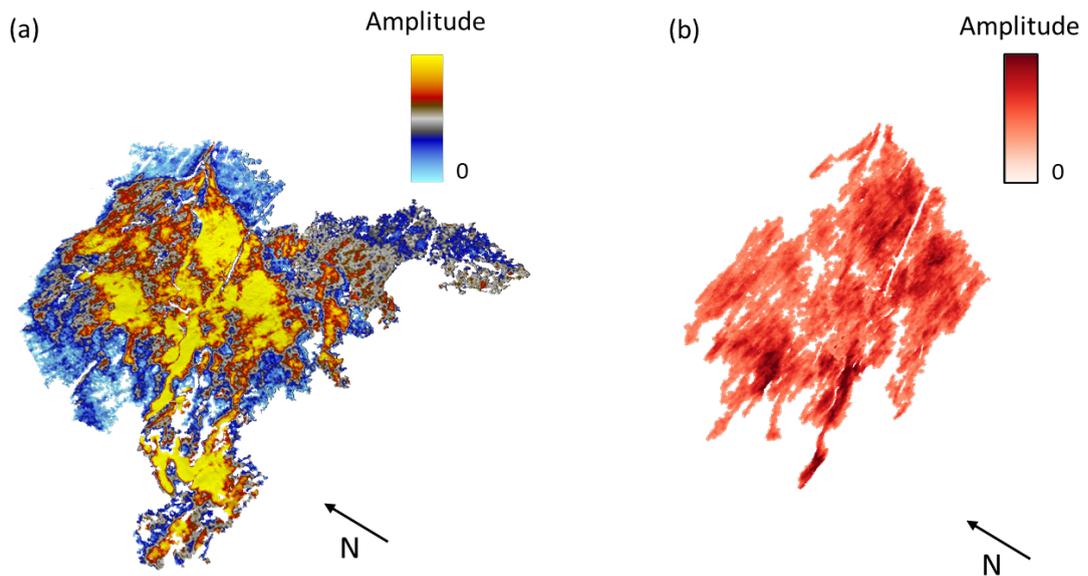


FIGURE 3.23: Comparison between the same turbidite fan extracted (a) by using Petrel seed-based auto-tracker, and (b) by using the point cloud segmentation method. The manual interpretation from Petrel is considered as ground truth for comparison.

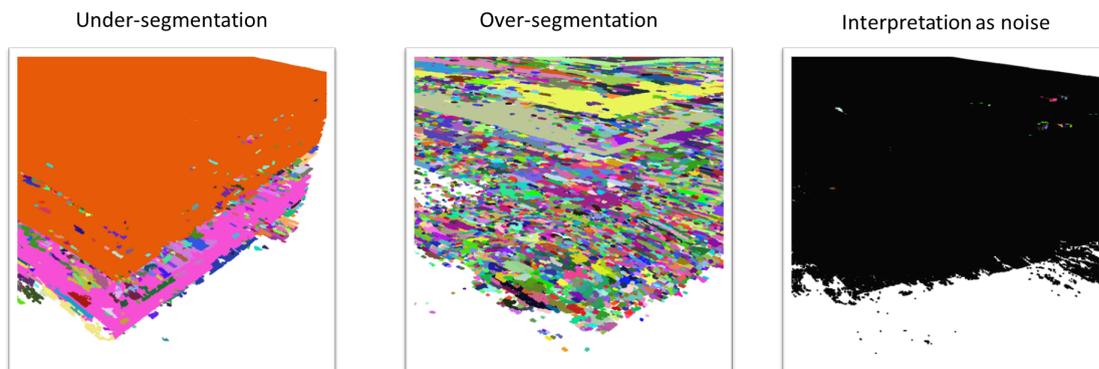


FIGURE 3.24: Three main segmentation problems have been identified in the segmentation of the seismic point cloud: under-segmentation, over-segmentation and interpretation mainly as noise.

3.3.4 Evaluation of the Segmentation

3.3.4.1 Methods

The ability to evaluate and reflect on the quality of the interpretation is fundamental. Especially with automated interpretation processes, evaluation becomes even more important to ensure some level of control. With the proposed seismic segmentation method, a result gets obtained in a matter of minutes. The result of the segmentation can be evaluated visually and qualitatively. However, it is not straightforward to make a decision visually, especially for seismic of large size. Another way to evaluate the quality of the segmentation can be indirectly by the effectiveness of the subsequent usage of the segmentation. For instance, we can base the evaluation on the success of the detection of particular geobodies. But, here, we want to provide a way to quantitatively assess the quality of the segmentation before doing any subsequent processing steps. We present some methods that provide evaluation metrics of the quality of the segmentation both at a global level and a cluster level.

Three main problems are identified in seismic segmentation: under-segmentation, over-segmentation and interpretation as noise for the better part of the seismic (figure 3.24). We often characterize under-segmentation as horizons merged together, but it could also refer to a more subtle feature segmented as part of a wider horizon. The over-segmentation often results in a very high proportion of tiny segments, making their

interpretation impossible. Generally, a balance always needs to be found in the segmentation between under-segmentation and over-segmentation. A segmentation mainly interpreted as noise is easy to identify. Still, it could lead to dramatic results if included in a fully automatic pipeline with no evaluation of the quality of the segmentation. These three problems could potentially lead to an incorrect segmentation of the geobodies in the seismic data and lead to an unsuccessful detection. Thus it justifies the need for a quantitative evaluation of the segmentation results early on.

Most of the segmentation evaluation methods present in the literature involve ground truth comparison. The accuracy is commonly evaluated compared to a segmentation manually labelled [Tarsha Kurdi and Awrangjeb, 2020]. However, a full segmented ground truth seismic is not available in this work. And from a practical point of view, a ground truth will not be available for real-world projects. Another approach is to evaluate the segmentation based on the homogeneity of the segmentation clusters or regions. The segmentation quality then derives from aggregating the local quality of its constituent segments. Zhang et al. [2003] propose a segmentation evaluation method based on information theory. They use pixel entropy to measure the pixel luminance (colour) uniformity within a region. Troya-Galvis et al. [2015] studied 5 homogeneity indexes: pixel entropy, contrast [Vojodi et al., 2013], CIELAB contrast [Chen and Wang, 2004], cohesion [Corcoran et al., 2010] and variance [Johnson and Xie, 2011] to quantify under- and over-segmentation of an image. They proposed two aggregation functions that one can use to combine these local quality estimates into a global quality score. Their experiment demonstrates that the suggested metrics can be utilized to identify under- and over-segmentation errors. However, the choice of the local quality estimates must be selected according to the type of segmentation task that one wishes to perform. Also, all these segmentation evaluation metrics were designed for image analysis. It needs to be adapted to the nature of point-cloud segments.

To evaluate the quality of the segmentation of the seismic point cloud, we propose to base the evaluation methods on intuitive quantitative criteria. Intuitively, a good segmentation would typically display:

- (i) A low number of noise points to avoid having only noise points classification from the segmentation.
- (ii) A large number of clusters of significant size. We expect a high number of significant clusters to be a good way to ensure a balance between under- and over-segmentation. Indeed, an under-segmented seismic would potentially lead to big segments (e.g. horizons) merged together. The number of big clusters is therefore reduced compared to what it would be if they were divided into individual big clusters. On the opposite, an over-segmented seismic would potentially lead to a big segment split into many small clusters, too small to qualify to be counted as clusters of significant size. Hence, again, reducing the number of significant clusters.
- (iii) Clusters would have a low ratio of superposed layers in the depth/time direction. By nature, the geological geobodies resulting from fluvial deposit phenomena are primarily laterally extended. Therefore, the overlap of two seismic events will generally indicate under-segmentation.
- (iv) Amplitude being homogeneously spatially distributed within a cluster. We expect continuity in the amplitude texture of a segment representing a single geobody. A hard edge contrast of amplitude would typically suggest that two different events are segmented together.

However, these four criteria should be analysed with flexibility regarding estimations of the segmentation quality. First, it is expected to have some noise remaining as the seismic reflectivity data is a noisy indirect measure. Second, it is expected that some geobodies might display an overlap. For instance, with turbidite fans, it is not unlikely to find some lobes overlapping or partially overlapping. Third, it is expected that some segments might display some hard edges of amplitude. For instance, the wide horizons are likely to contain some minor faults that might have affected the amplitude distribution and provoked some hard edges in the amplitude distribution.

All these four criteria translate into four metrics used to evaluate the segmentation. The first two metrics are straightforward. The metric considered for assessing the proportion of noise points is simply the ratio of noise point (r_{noise}) in the seismic (equation 3.3).

$$r_{noise} = \frac{n_{noise_points}}{n_{points}} \quad (3.3)$$

Regarding the number of significant clusters, the interpreter inputs a minimal number of points ($\delta_{big_clusters}$) to the evaluation. We consider a cluster with more points than this threshold as significant. The choice of this number of points should depend on the seismic size and resolution. In order to keep a score between [0;1], the ratio of significant clusters is computed ($r_{big_clusters}$). We obtain this ratio by dividing the number of big clusters by the maximum possible number of big clusters in the seismic (equation 3.4). The maximum possible number of big clusters in the seismic point cloud is the integer portion of the division of the number of points in the point cloud by the threshold defining the minimal size of a big cluster.

$$r_{big_clusters} = \frac{n_{big_clusters}}{int(n_{points}/\delta_{big_clusters})} \quad (3.4)$$

We estimate the ratio of overlapping layers for each cluster. We compute an approximation by counting the number of overlapping points in the trace direction in the cluster, the Horizontally Superposed Layers index - *HSL* - (equation 3.5). For each unique (x,y) point coordinate in the cluster, we sum the number of points with this given coordinate. Then, we divide this number by the number of unique pairs to obtain a ratio. Figure 3.25 illustrates the count of overlapping points on two traces.

$$HSL(R_j) = \frac{1}{S_{(x,y)_j}} \sum_{x \in X_j} \sum_{y \in Y_j} N_{(x,y)_j} \quad (3.5)$$

where X_j is the unique coordinate x in the region j (R_j); Y_j is the unique coordinate y in R_j ; $S_{(x,y)_j}$ is the number of unique pair of coordinates (x,y) in R_j ; and $N_{(x,y)_j}$ is the number of points with coordinates (x, y) in R_j . We compute the mean of the *HSL* of

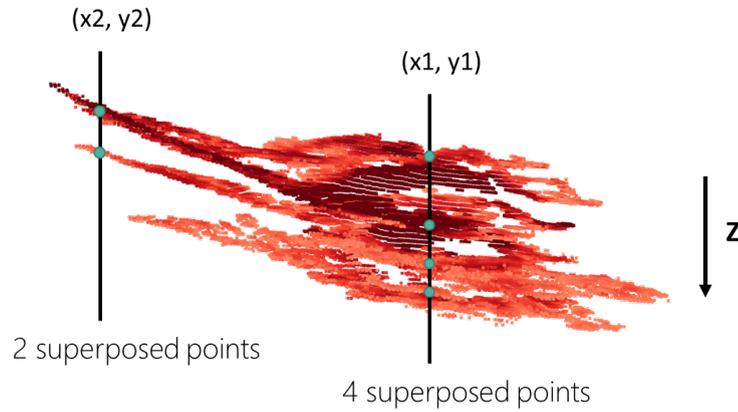


FIGURE 3.25: The point cloud is a segment of the seismic reflectivity data displaying overlapping structures. The count of overlapping points is illustrated for two doubles of (x, y) coordinates.

all clusters to obtain a global metric. To get a more representative score, we weigh the mean of HSL by the number of points in the cluster ($\overline{HSL_w}$).

For the estimation of amplitude homogeneity within each segment, we first tried the region entropy [Zhang et al., 2003], adapted to a point cloud (equation 3.6).

$$H(R_j) = - \sum_{m \in V_j} \frac{L_j(m)}{S_j} \log \frac{L_j(m)}{S_j} \quad (3.6)$$

where V_j is the set of all possible values in R_j , $L_j(m)$ is the number of points in R_j that have the value m , and S_j is the number of points in R_j . This definition of entropy is designed for analysing the distribution of a categorical variable. To conform to this definition, we first classify the amplitude distribution in each segment into 100 bins. Hence, V_j represents these 100 possible categories. And $\frac{L_j(m)}{S_j}$ can be seen as the probability of a point belonging to a specific category. A better solution would be to use the limiting density of discrete points formula [Jaynes, 1968] which extends the idea of Shannon entropy [Shannon, 1948] to continuous variables.

However, the entropy provides a measure of the overall amplitude distribution of a segment, but it does not measure its spatial distribution. The spatial distribution of amplitude is essential to determine the consistency of a segment. Indeed the continuity

of amplitude often reflects some depositional patterns and thus are indicators of points belonging to the same depositional event.

We propose to use Zeboudj's internal disparity (CI) adapted to measure the spatial distribution of amplitude in a point cloud (equation 3.7). CI was first introduced by Zeboudj [1988] and used by Vojodi et al. [2013] as an evaluation criterion of the internal disparity of a property in a region.

$$CI(R_j) = \frac{1}{S_j} \sum_{s \in R_j} \max(C(s, t), t \in W_\delta(s) \cap R_j) \quad (3.7)$$

with:

$$C(s, t) = \frac{Amp(s) - Amp(t)}{Amp_{max}} \quad (3.8)$$

where S_j is the number of points in the region R_j . $W_\delta(s)$ are the ensemble of neighbouring points at a distance δ of the point s . δ is a parameter that needs to be selected when performing the evaluation. It is the sphere's radius centred on s defining its neighbourhood. We should choose δ considering the maximal distance between two points that should not differ too much to have a homogeneous spatial amplitude distribution. It should typically be a low number such as 3 for a seismic dataset, where point resolution is 1 in the three dimensions. $Amp(s)$ is the amplitude value of point s . Amp_{max} is the maximum amplitude value in the seismic reflectivity data. In simpler terms, $CI(R_j)$ can be described as the mean of the maximum amplitude difference between each point and its immediate neighbouring points of a segment. Figure 3.26 presents the results of entropy value and Zeboudj's internal disparity of two point clouds. The first one (figure 3.26.a) is homogeneous in amplitude while the second one (figure 3.26.b) has a more heterogeneous spatial distribution. This difference between the two clusters is well captured by $CI(R_j)$ – two times higher in the second point cloud. But it is less reflected in $H(R_j)$ - only a 5% difference of $H(R_j)$. This is due to $H(R_j)$ being based on the discretized amplitude distribution (figure 3.26 c & d) and not reflecting the spatial distribution of the amplitude. The efficiency of $CI(R_j)$ versus $H(R_j)$ to distinct homogeneous and heterogeneous regions are illustrated in Figure 3.27. It shows for each metric 8 point clouds corresponding to 4 highs and 4 lows sampled among

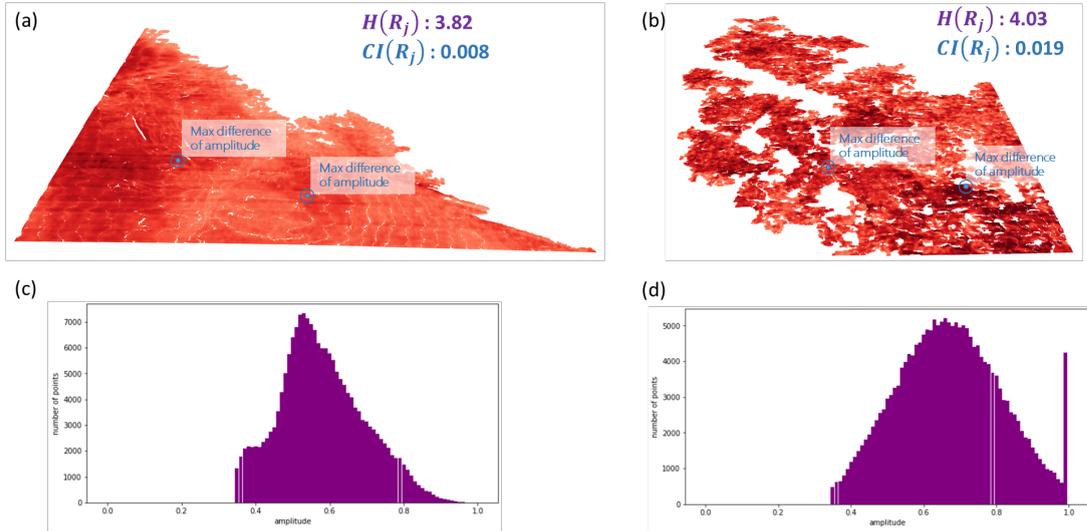


FIGURE 3.26: Illustration of the amplitude entropy (purple) and the Zeboudj's internal disparity (blue) on an homogeneous point cloud (a) and an heterogeneous point cloud (b). (c) and (d) display the discretized amplitude distribution of the point clouds, respectively (a) and (b). That discretized amplitude distribution is used to compute the amplitude entropy.

the clusters of the *validation dataset*. For $CI(R_j)$, we observe a clear distinction of amplitude homogeneity. However, we cannot make such a clear distinction for the 8 point clouds sampled for $H(R_j)$.

Likewise, the mean of the $CI(R_j)$ of all clusters is computed, weighted by the number of points in the cluster, to provide a global metric for the entire segmentation result.

We can use all these four metrics to evaluate a segmentation result of a seismic cube. The amplitude continuity approximated by CI is theoretically a good approach to evaluate cluster consistency. However, practically, the seismic amplitude is too noisy to provide a good indicator of the segmentation quality. Hence we propose not to use it for the global quality score of the segmentation (β). We propose to combine the tree metrics r_{noise} , $r_{big-clusters}$ and \overline{HSL}_w as in equation 3.9 to obtain a global segmentation quality score ($\beta_{segmentation}$).

$$\beta_{segmentation} = \frac{r_{big-clusters}(1 - r_{noise})}{\overline{HSL}_w} \quad (3.9)$$

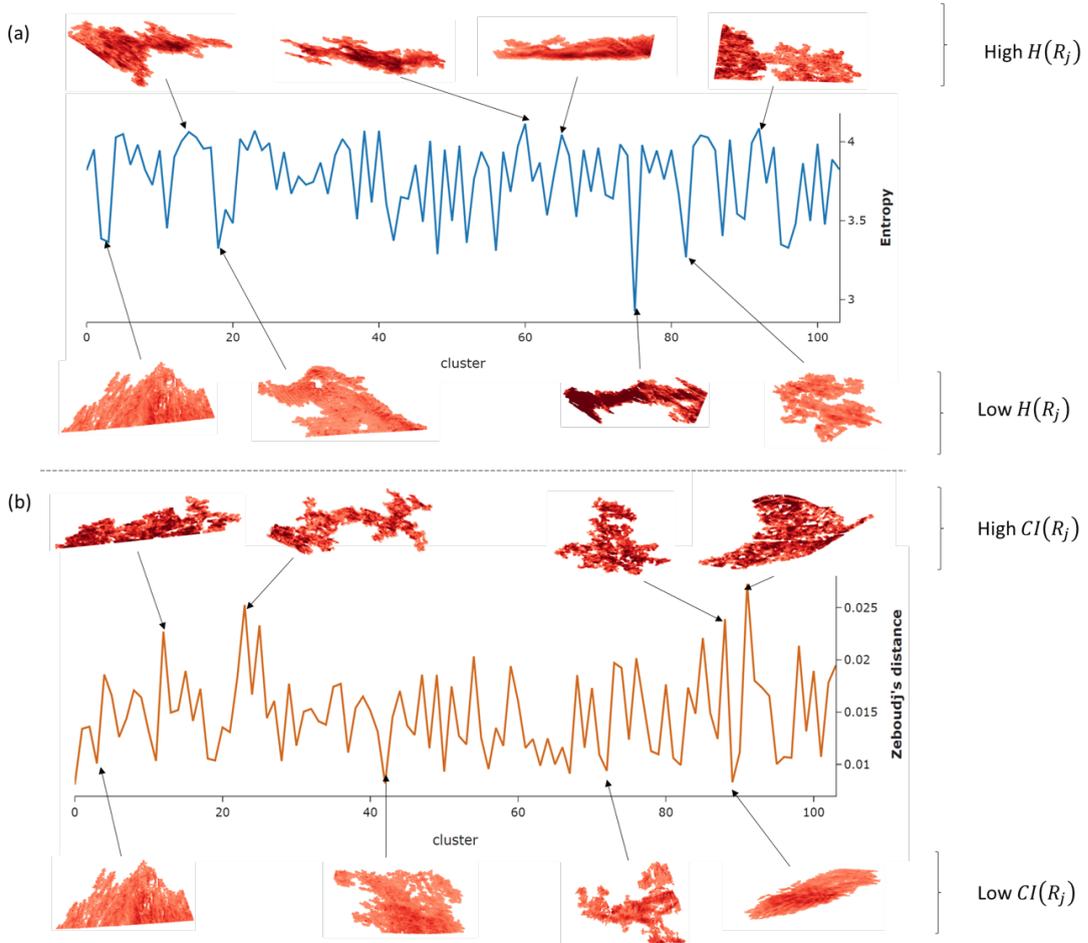


FIGURE 3.27: For each metric Entropy (a) and Zeboudj's internal disparity (b), 4 high metric value point clouds and 4 low metric value point clouds are sampled and displayed among all the development test clusters.

However, $\beta_{segmentation}$ should not be interpreted as an absolute metric for evaluating the quality of the segmentation but rather as a relative metric to compare two different segmentation results and evaluate their relative quality.

In addition to this global segmentation score, HSL and CI provide a cluster-level evaluation. $HSL(R_j)$ give an approximation of the overlapping structures of the cluster. And $CI(R_j)$ give an approximation of the homogeneity of the cluster. These two metrics can help sort the clusters based on the quality of their segmentation. Thus it allows for identifying clusters with segmentation problems which might need further optimisation.

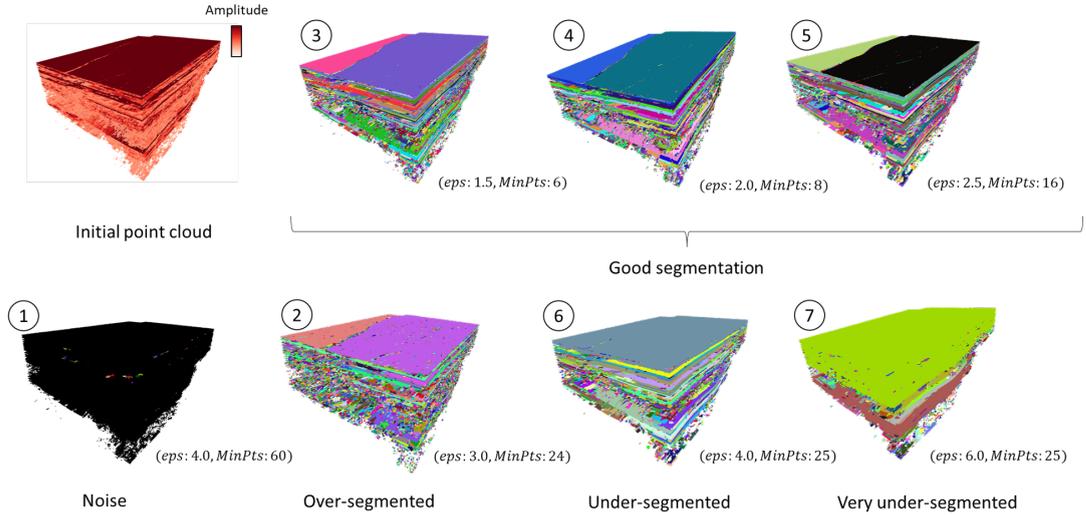


FIGURE 3.28: Validation set for the evaluation of the segmentation. This 7 realisations of the segmentation of the *validation dataset* were obtained by varying the DBSCAN parameters for the segmentation (*eps* and *MinPts*). The realisations are sorted from over-segmented to under-segmented and identified by a number between 1 and 7.

Evaluation of the 7 segmentation realisations

Realisation	1	2	3	4	5	6	7
r_{noise}	0.9950	0.5751	0.2022	0.2396	0.2067	0.1496	0.0390
$r_{big-clusters}$	0.0025	0.0374	0.0448	0.0468	0.0448	0.0353	0.0064
HSL_w	1.7390	1.0880	1.2450	1.1370	1.2618	2.5632	16.5107
$\beta_{segmentation}$	0.000007	0.014599	0.028704	0.031270	0.028160	0.011712	0.000371

TABLE 3.1: Global segmentation scores of the 7 realisations obtained from the *validation dataset*.

3.3.4.2 Validation on the *validation dataset*

Global segmentation score

In order to validate this global segmentation score, we propose the following experiment: 7 different realisations of the segmentation of the *validation dataset* were realised by choosing different *eps* and *MinPts* values (figure 3.28). These the 7 realisations were analysed and sorted from over-segmented to under-segmented. We judged one realisation as mainly noisy, one as over-segmented, two from under-segmented to very under-segmented, and three realisations as good segmentation.

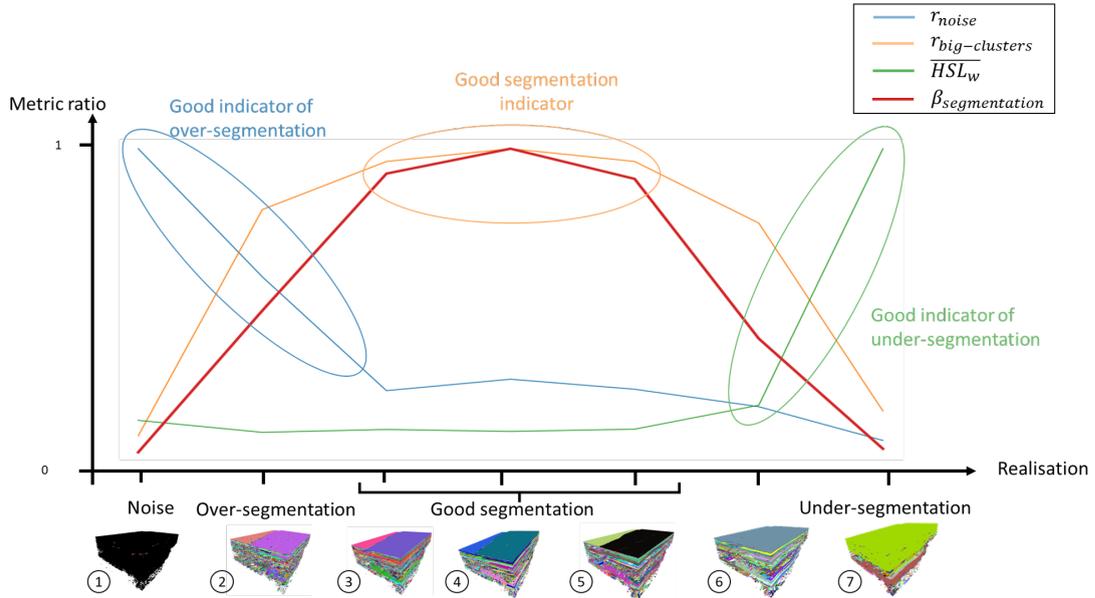


FIGURE 3.29: Variation of the global segmentation metrics versus the different segmentation realisations sorted from over-segmented to under-segmented. To be visible on the interval $[0; 1]$, each metric is plotted as a ratio over the maximum metric score of that metric for the 7 realisations.

We present the global segmentation scores of each realisation in Table 3.1. Figure 3.29 presents the variations of these global segmentation metrics against the different segmentation realisations sorted from over-segmented to under-segmented. r_{noise} turns out to be a good indicator of over-segmentation, \overline{HSL}_w is a good indicator of under-segmentation, and $r_{big-clusters}$ is a good indicator of a good segmentation. $\beta_{segmentation}$ judged the 3 "good-segmentation" realisations with significantly higher scores. The results suggest that the realisation 4 presents the best segmentation quality. Amongst the seven values sampled pairs, the best segmentation parameter values are an eps of 2.0 and a $MinPts$ of 8. Overall, $\beta_{segmentation}$ seems to be an appropriate metric to evaluate the quality of the segmentation of a seismic data point cloud. This experiment validates the use of $\beta_{segmentation}$ to compare relative segmentation qualities of seismic data point clouds.

Cluster-level segmentation score

We also evaluated the segmentation quality at a cluster level. HSL and CI provide a score of segment quality for each cluster. Thus, HSL and CI allow sorting the clusters by quality. Figure 3.30 presents the results of both metrics for each cluster with more than

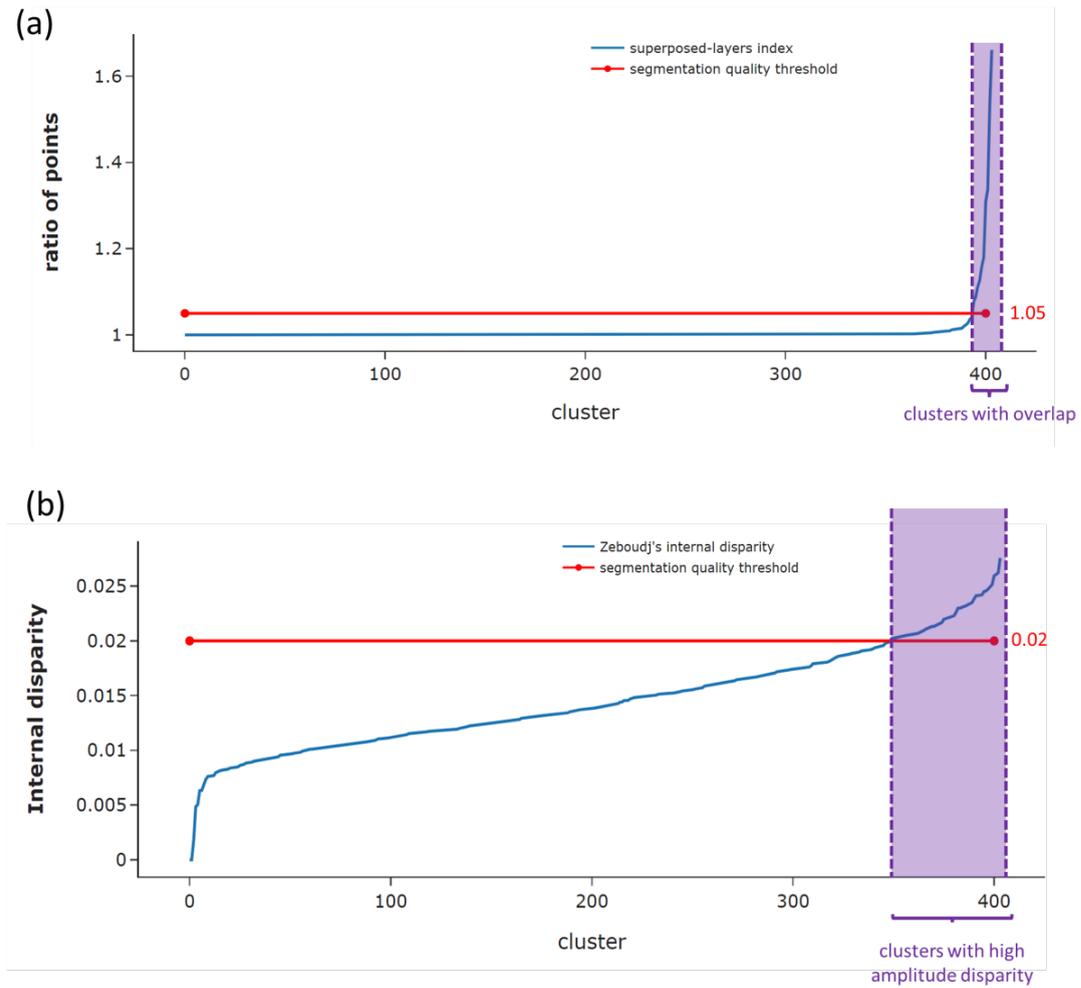


FIGURE 3.30: Segmentation evaluation metric at a cluster-level - (a) *HSL* and (b) *CI* - for the 404 main segments of the *validation dataset*. In red, we plot a quality threshold for each metric to identify segments with high overlap and high amplitude disparity.

5,000 points from the *validation dataset*. We note little overlaps in most of the clusters. Only 10 clusters have more than 5% of their points overlapping. *CI* was computed on a neighbourhood radius of 2.5. The internal disparity of the clusters is low and evenly distributed over all clusters. We note 55 clusters possess an amplitude disparity higher than 0.02. The metrics used at a cluster level offer a solution to identify point clouds with overlap or heterogeneous amplitude distribution quickly.

3.3.5 Optimisation of the Segmentation

3.3.5.1 Methods

After an initial segmentation of the seismic reflectivity data, we can improve upon that first segmentation. Indeed, achieving a perfect full-volume seismic segmentation for all structures in the cube is not realistic. Seismic reflectivity data are often too noisy, and the structures are too varied for a single set of DBSCAN parameters to produce a perfect segmentation. We propose to improve the initial segmentation by identifying bad quality segments and computing a new segmentation of these clusters.

There are two scenarios for optimising a bad quality segment: (i) it can be merged with neighbouring segments – solving over-segmentation errors, and (ii) it can be divided into more uniform segments – solving under-segmentation errors.

Regarding the first scenario, *merge operations*, we propose to base the merge of neighbouring segments on the condition that the union of the two segments as one segment is better than their division. Figure 3.31 presents the conceptual flowchart for optimising an initial segmentation through the merge operations. First, we input the initial segmentation obtained with DBSCAN. Then, we iterate over the segments and evaluate the possible merge with a neighbouring segment. We define neighbours as the segments at a distance less than the value input for optimisation. In addition, the two segments assessed must have a better segmentation score considering the metrics presented in section 3.3.3. If so, we merge the two segments and assess the next segment until all segments have been assessed. However, in practice, this iterative process is very time-consuming, and its application might reduce the performance of the seismic data segmentation method presented here. Moreover, the benefit is limited compared to the optimisation provided. Thus, we will not apply this optimisation for the main seismic applications presented in chapter 5.

Regarding the second scenario, *division operations*, we propose to base the division on a finer segmentation performed only on the segments that require further optimisation. We carry out several segmentations with DBSCAN set up with different hyper-parameters

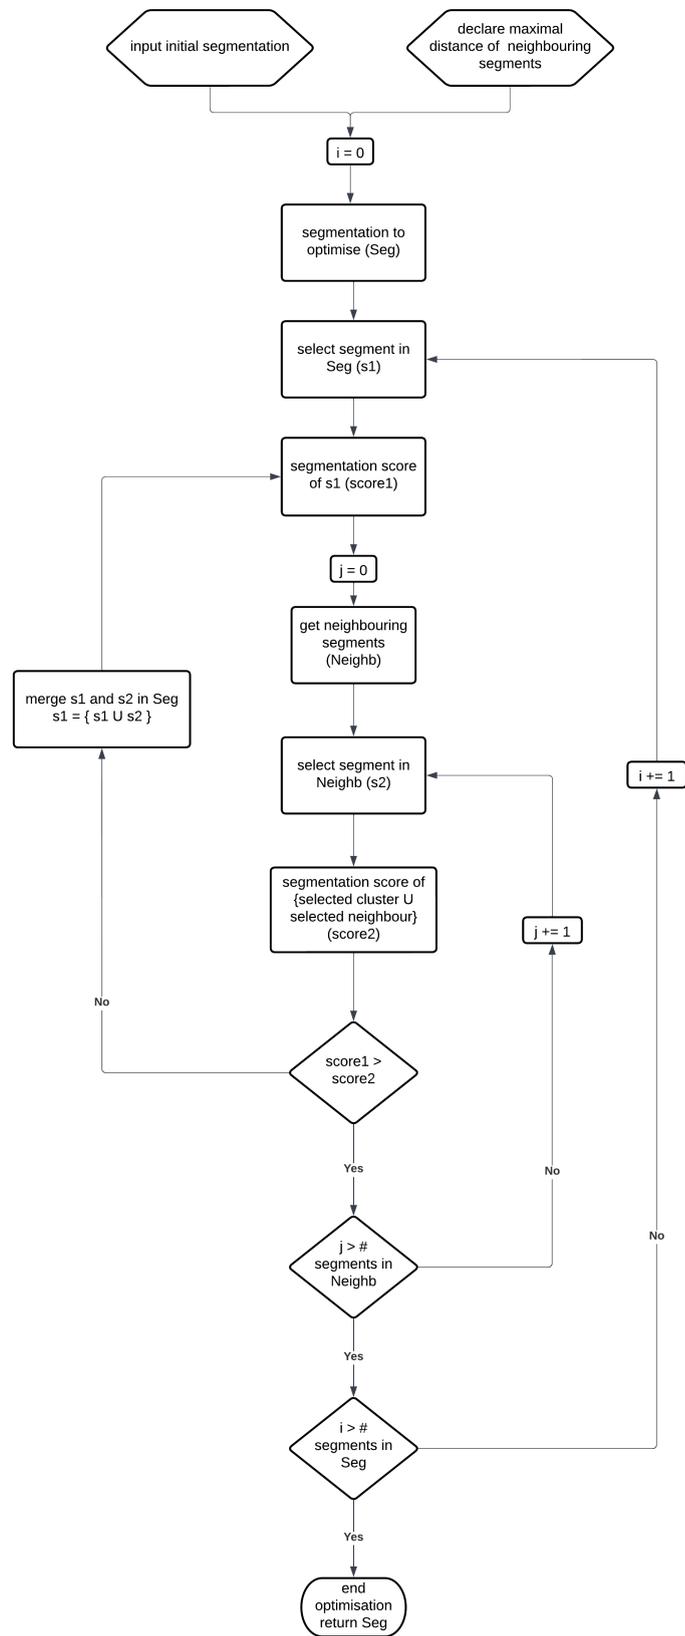


FIGURE 3.31: Flowchart of the optimisation of the segmentation by merge of neighbouring segments, *merge operations*.

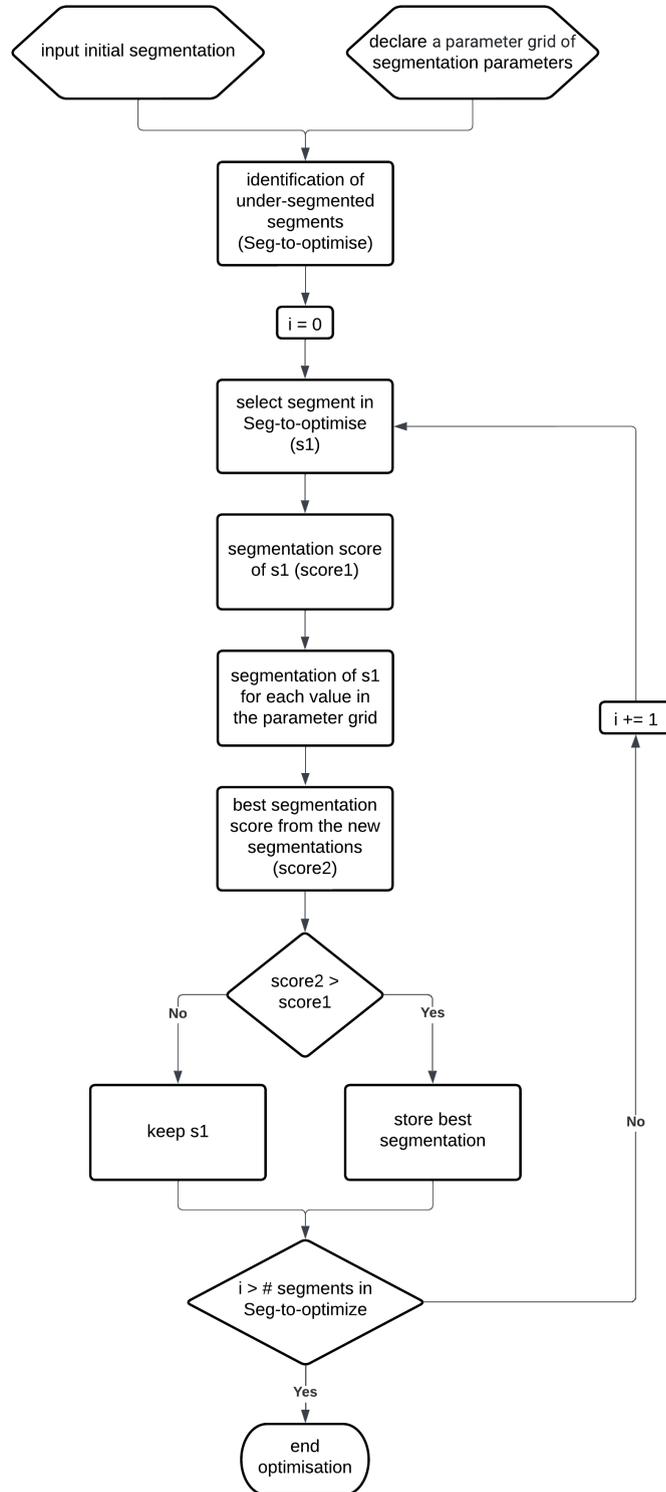


FIGURE 3.32: Flowchart of the optimisation of the segmentation by division of under-segmented segments, *division operations*.

and keep the best realisation. Figure 3.32 presents the conceptual flowchart for optimising an initial segmentation through the division operations. Our first step consists of identifying the segments of bad segmentation quality. As presented in section 3.3.3, *HSL* is a good metric to identify under-segmented clusters. To further refine this selection, we can perform a visual validation of the selected segments to narrow it down to the essential segments. Then, we define an exhaustive search over DBSCAN segmentation parameter values to find the best parameter set. The parameter space to explore must be input at the start of the optimisation process. Again, we get the best segmentation considering $\beta_{segmentation}$. We keep the division if its score is higher than the score of the initial segment. And we repeat this process for all segments with an under-segmentation.

3.3.5.2 Validation of the optimisation

In order to validate the *division-operation* optimisation, we apply that procedure to an example segment. We selected the example segment because it had a high *HSL* value and presented interesting geological structures. Indeed, two turbidite fans are clustered in the same cluster. We can tell there are two different fans because the sediments came from two distinct feeder systems (figure 3.33). The validation procedure consists of testing if the *division-operation* optimisation algorithm can distinguish the two clusters.

The optimisation successfully extracted a segmentation differentiating the two turbidite fans. Figure 3.34 shows the results of the optimisations. We vary *eps* between 2.2 and 2.95 with an increment of 0.15, and we vary *MinPts* between 14 and 23 with an increment of 3, which leads to a total of 24 realisations. For each realisation we compute $\beta_{segmentation}$. The total processing time for the optimisation process is 19s for a cluster of 242,150 points. We observe that 12 realisations differentiate the two turbidite fans. And the 6 best scores accurately distinguish the two fans. The small differences in scores are due to noise points and very small clusters interpreted differently between the realisations.

A more sensible choice of the parameter values can improve the grid search. Indeed, we observe eight realisations segmented only as noise in this example. It is due to the choices of pairs of (*eps*, *MinPts*) leading to a density of points too high to find any core

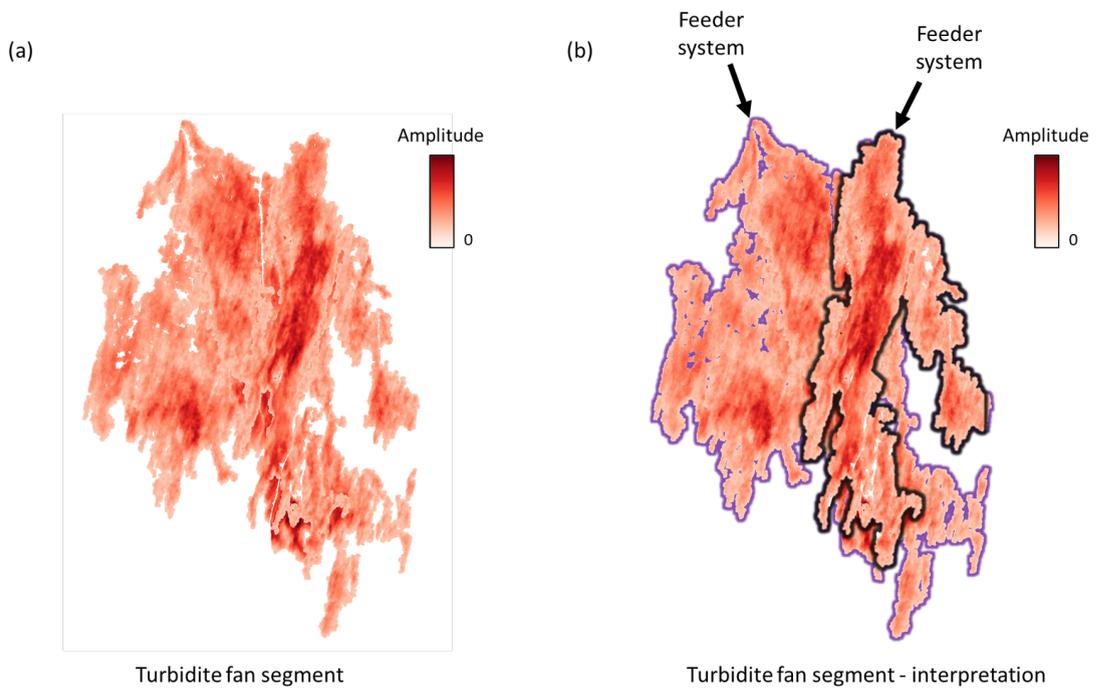


FIGURE 3.33: Under-segmentation problem of two turbidite fans being segmented in the same cluster.

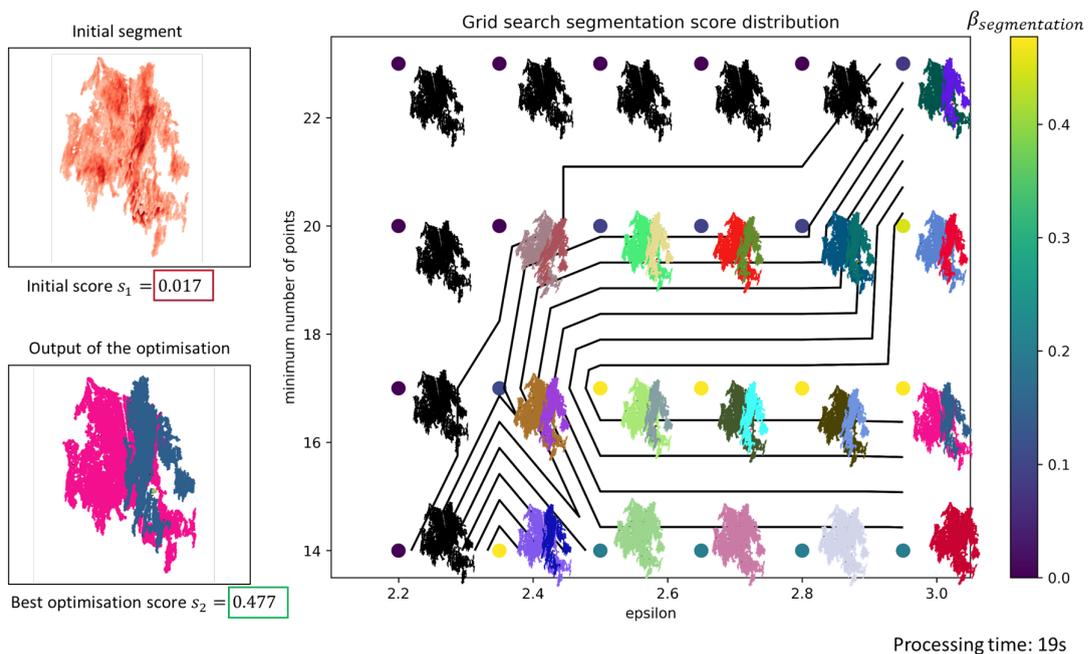


FIGURE 3.34: Grid-search over ϵ and $MinPts$ and corresponding segmentations. The optimal realisation segments the initial segment in two main clusters differentiating the two turbidite present in this initial segment.

point for this point cloud. Analysing the point density beforehand would lead to these realisations being disregarded, saving time in the optimisation process.

3.4 Conclusion and discussion

In this chapter, we presented a fast and robust workflow to perform a segmentation of the full seismic volume at once. The workflow starts with filtering of the seismic data through extrema extraction, weak amplitude cut-off and low coherence cut-off (but other seismic attributes could be used). The filtering produces a point cloud representation of the seismic data. The point cloud is then segmented into spatially connected clusters using the DBSCAN algorithm. Every step of the workflow outputs an interactive visual, allowing the choice of parameter values to be adjusted. As a result of the workflow, a seismic cube is collapsed into segmented objects that can be quality controlled and interpreted towards the target geological objects.

Furthermore, we presented methods to quantify the quality of the obtained segmentation and improve it where necessary. The quality evaluation is based on predefined metrics analysing the segmentation both at a global level and a cluster level. Manual quality checks are still required to better assess the realism of the segmentation, in particular when additional geological information is known. However, these quantitative metrics provide a quick and unbiased score to a segmentation realisation and detect poor quality segments. We perform a parameter search on DBSCAN hyper-parameters to improve the poor quality segments and find a more optimal realisation. Not discussed here and that could be done, there is the possibility to include the amplitude cut-off value and coherence cut-off value as parameters to explore in the optimisation process.

The proposed workflow meets all the requirements outlined in the introduction of this chapter. It segments the seismic into connected objects regardless of their scales: extended horizons and more subtle events such as fan deposits. It is extremely fast to compute. Working with a point cloud format makes the seismic easier to manage and quicker to process. Indeed, the point cloud extraction reduces the volume of data by a factor of 100. Moreover, treating the segmentation as a density-connectivity problem

with DBSCAN decreases the computational complexity compared to treating it as an optimisation or region-growing problem.

The seismic segmentation performed here does not aim to replace the existing manual expert seismic interpretation workflow. Instead, it provides a new perspective on how to look at seismic data with respect to interpretational uncertainty. The workflow enables a reduction in routine horizon picking and allows to focus on interpretational uncertainty. The segmentation outputs the main horizons. However, the horizons here are clusters of points, not surfaces. For sequence stratigraphic seismic interpretation, to construct relative geological age, the interpreters often need to work with a stack of surface-horizons that are smoother than the point-horizons directly extracted from the seismic data [de Groot et al., 2010, Paumard et al., 2019]. Extracting surface-horizons includes an additional degree of geological interpretations, moving from the data toward a model representation.

The primary purpose of applying this seismic segmentation method is geobody detection, but it also benefits well correlation frameworks. The sparse visualisation offered by the point cloud representation allows for visualising the continuity of geological events between the wells. Moreover, our seismic segmentation provides a level of detail higher than the standard seismic interpretation workflows. The point cloud segmentation provides the main stratigraphic horizons constraining the main correlations between wells and also the subtle bodies, which allow a refined comprehension of event continuity across wells -geobodies correlatable or not-.

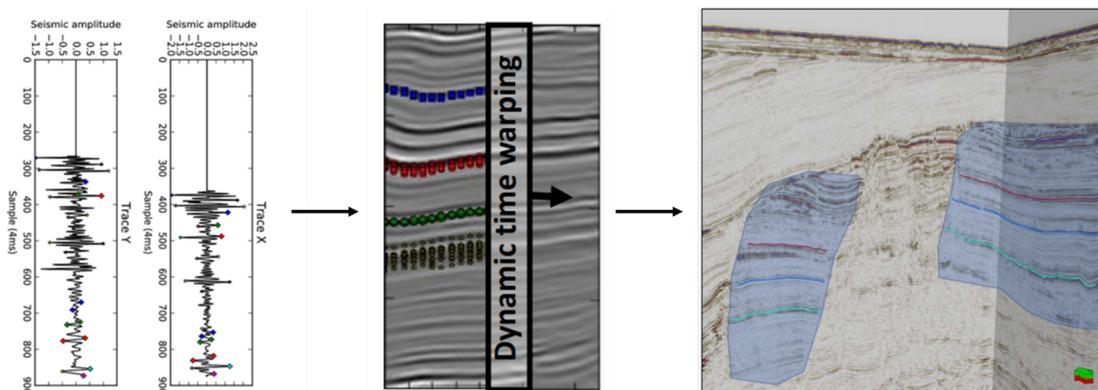


FIGURE 3.35: Horizons tracked across fault blocks using Dynamic Time Wrapping for signal correlation. Figure modified from [Bugge et al., 2019].

When using this seismic segmentation method, a few limitations and drawbacks are important to keep in mind. Firstly, filtering on coherence leads to a loss of information at the fault areas. So, we might be missing features located in heavily faulted regions. Secondly, if the signal-to-noise ratio in the seismic data is low, it will impact the quality of the segmentation. Indeed, a high proportion of noise can lead to two undesired outcomes: (i) the loss of lots of information at the coherence filtering or (ii) an under-segmentation of the seismic due to noise points creating false connectivities between seismic events. Finally, the segmentation does not include across-fault interpretations. In the case where an event is faulted from part to part in the seismic, the segmentation will not find connectivity between the points. Instead, it will interpret it as two distinct clusters. Adding across fault block horizon tracking as a post-processing step to the segmentation could be one approach to resolve this problem. For instance, the dynamic time warping (DTW) algorithm [Müller, 2007, Sakoe and Chiba, 1978] can find matching sequences in two seismic traces across two fault blocks [Bugge et al., 2019, 2018, Jin et al., 2017] (figure 3.35). DTW matches effectively and accurately one-dimension sequences or time series with non-linear oscillations along the time axis. Moreover, it is robust to shifts and missing sequences, which is essential when comparing seismic traces [Hale, 2013].

Chapter 4

Detection of 3D geobodies in a segmented seismic

4.1 Introduction

This chapter introduces a new workflow for detecting specific geobodies of interest within seismic volumes from the segmented point cloud data. Object recognition and retrieval constitute the second part of the methodology that directly follows the seismic segmentation presented in chapter 3 (figure 1.7). The seismic segmentation provides a delineation in 3D of all objects present in a seismic volume. The idea of geobody detection is to add further semantic information to the segmented objects and enhance their geological interpretation. Thus, the objects will no longer be just clusters of points but will gain geological meaning with the elicited interpretation attached to them.

We work on the simpler problem of finding items that look like a geobody of interest and screening them rather than classifying the entire volume for all geobody types. An ideal solution would be to label every object with a geological type or meaning. However, this is a highly complex and subjective task due to the diversity of geological labels and object types possibly present in the subsurface. Therefore, here, we reduce the problem of detecting a specific type of geobody amongst all objects in the seismic data. However,

the proposed methodology is a general purpose and does not depend on the specifics of the target geobody to detect.

4.1.1 Seismic attributes for geobody detection

The current most common approach to detect geobodies uses a combination of seismic attributes to highlight and delineate the geobodies in the seismic. The most popular attribute is spectral decomposition. First proposed by Partyka et al. [1999], spectral decomposition is a powerful aid to imaging thin-bed reflection by exploring the data in the frequency domain [Marfurt and Kirlin, 2001]. Thanks to multi-channel RGB visualisation, spectral decomposition can capture three frequency bandwidths as three colour channels to help identify subtle geological features not discernible with traditional seismic visualisation techniques [Henderson et al., 2007, Lawal et al., 2022, Purves et al., 2007]. Thus, spectral decomposition has been successfully used to support the interpreter to better image channels and faults (figure 4.1). Other attributes amplitude-derived attributes such as coherency, root-mean-square amplitude (RMS), signal envelope or sweetness [Chopra and Marfurt, 2005, Hart, 2008] are also helpful in detecting and delineating channels and other stratigraphic features manually. The results of these analyses improve the geometry and delineation of geobodies and their inner facies fill [Monier et al., 2021].

Seismic attribute image analysis does not solve the problem of automated detection of geobody. Automated detection implies an algorithm to output the position and type of geobodies, while seismic attributes allow a better definition of a geobody before manual detection can be performed. These different seismic attributes enhance the physical property contrast -e.g. lithological or fluid variations- on a seismic slice, allowing the interpreter to delineate geobodies better. However, producing a good mapping of these seismic attributes often requires prior knowledge of the position of a particular geobody in order to extract a slice between the top and base of the geobody that the interpreter wishes to map and to work with an accurate window of interest. Thus, this chapter aims to go one step further towards automatically detecting specific geobodies based on their characteristic geometries, shapes and seismic attribute distributions.

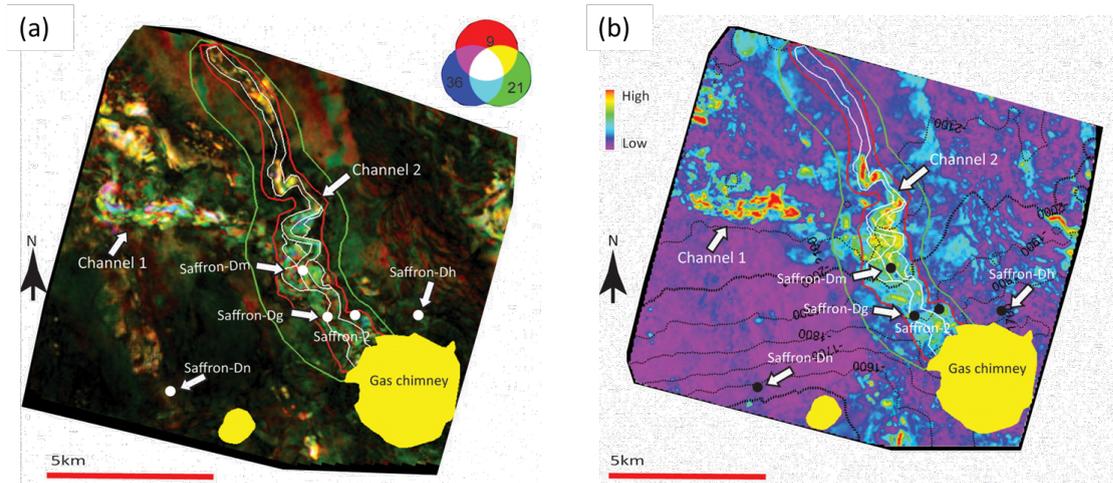


FIGURE 4.1: A channel manually delineated shown on (a) spectral decomposition RGB colour-blended map and (b) RMS amplitude extraction map. Both maps were extracted between the top and base channel grids. Modified from Monier et al. [2021].

4.1.2 Deep learning approaches applied to geobody detection

4.1.2.1 State of the art results

Recently, several successful attempts at automatic geobody detection with deep learning have been made.

Deep neural networks first delivered successful results on salt diapir and fault detection by learning to classify and predict a mask on the desired geobody. Salts and faults fit well with this image classification solution because they can be detected and characterised from 2D images [Di et al., 2018, Shi et al., 2019, Waldeland and Solberg, 2017, Waldeland et al., 2018, Wu et al., 2019, Xiong et al., 2018]. Indeed salts have poor internal reflectivity, leading to a broken, noisy pattern of reflectors in the seismic reflection [Jones and Davison, 2014] identifiable in the section. In the same way, faults cause discontinuities in reflectors observable in 2D.

Pham et al. [2019] go one step further in deep neural network solutions with the detection of channels, which are more subtle seismic features, and harder to spot and label on section views. Convolutional neural networks (CNN) trained on 3D samples demonstrated promising results for automatic interpretation and quantitative analysis of channels [Pham and Fomel, 2020, Pham et al., 2019]. However, a limitation of using

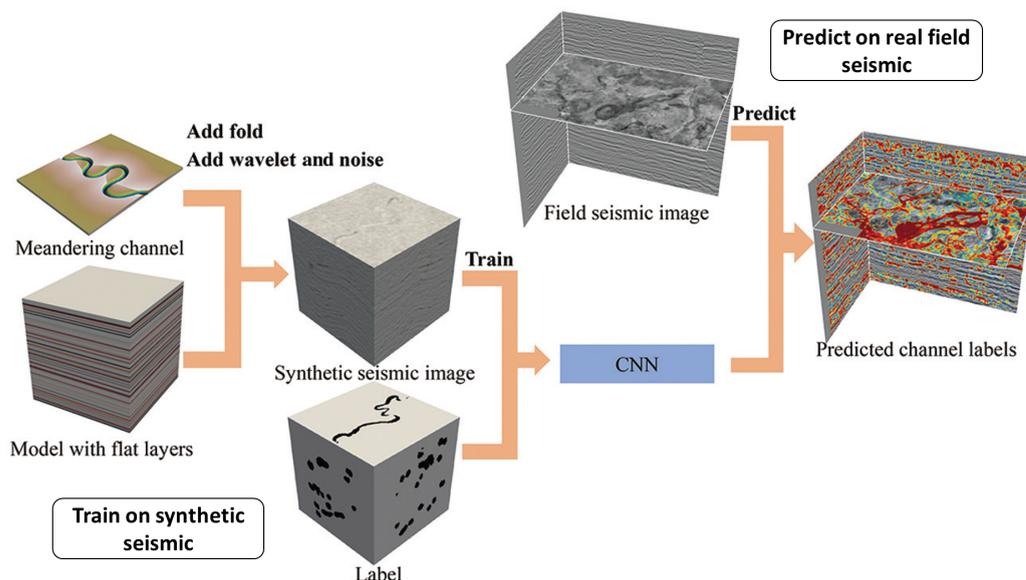


FIGURE 4.2: Flowchart of training a CNN for channel detection on synthetic seismic before applying to field seismic. Modified from Gao et al. [2021].

a CNN to detect channels is the need for a seismic dataset where all channels are labelled for training, something Gao et al. [2021] overcame by using a CNN trained on synthetic data before applying it to a real seismic cube (figure 4.2). Gao et al. generated the synthetic seismic by modelling realistic channel shapes using geologic numerical simulations (python *meanderpy* by Sylvester et al. [2019]). They also added structural deformations and noise to the training seismic to increase realism. Applied to two real seismic volumes, the 3D CNN proved efficient in detecting channels despite being trained on synthetic data. However, Gao et al. [2021] also put forward a limit of the method linked to the dependence on the diversity of the training data, mentioning that the method struggles with large-scale incised channels because these structures were not present in the training set. Indeed, if one wants to detect channels in general, having all types of channels represented in the training set is better. Deep learning often performs poorly to generalise on samples out of the range of representation present in the training set.

4.1.2.2 Challenges of deep learning applied to turbidite fan detection

The trained CNN approach is hard to adapt to turbidite fan detection due to the difficulties in generating representative labelled data for training, the size of turbidite fans and the

diversity of fan characteristics (figure 1.5). The lack of available, good quality 3D seismic data displaying interpreted turbidite fans makes it complex to create a labelled training dataset with enough examples to train a neural network. Moreover, labelling fans is a complex and subjective process. So, having to label an entire seismic cube for fans would likely lead to a significant portion of false-positive and false-negative labels in the training set.

Using a synthetic seismic could resolve the lack of available labelled training data but modelling realistic fan geobodies is an issue. By combining stratigraphic forward modelling of turbidites, which creates geological models by simulating dynamic sedimentary processes, with seismic forward modelling, one could produce synthetic seismic profiles to train a neural network [Wan et al., 2022]. Burgess et al. developed a reduced-complexity model that produces 3D representations of turbidite fan strata that would be well suited to creating synthetic seismic. [Burgess et al., 2019] - (figure 4.3). However, these reduced-complexity models are highly idealised explanatory models and do not reproduce the complexity and variety of turbidite fans present in real-field seismic [Bokulich, 2013].

Another potential issue with using CNN for turbidite fan classification is the sample size input in the neural network. Turbidite fans have incredibly variable shapes and dimensions, from 5 to >30 km in width, 5 to >140 km in length [Garcia et al., 2015]. So, inputting a small 3D cube in the neural network will result in only a cropped part of the turbidite fan being input. However, knowing how important the channel feeder and fan edges are for the accurate detection of a fan, having only part of the fan characteristics in the training samples could lead to biased learning of fan classification.

Although the deep learning workflow with a trained CNN has merits and deserves further investigation, in this chapter, we decided to explore another avenue to the problem of object detection in the seismic.

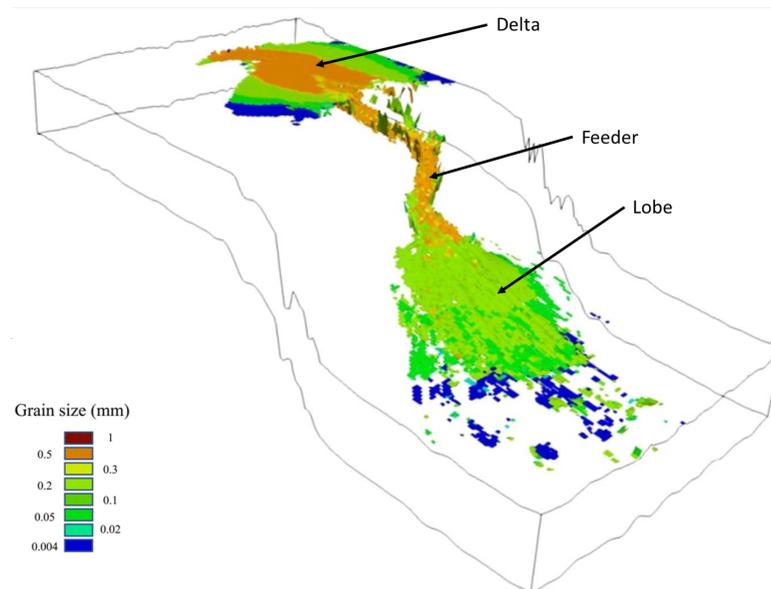


FIGURE 4.3: Turbidite fan deposit modelled with stratigraphic forward modelling. Modified from Wan et al. [2022]

4.2 Overview of our approach

In our approach, we aim to take advantage of the seismic segmentation provided by the methods introduced in chapter 3 to work directly on the 3D point cloud objects to characterise them and identify specific geobodies.

4.2.1 Object recognition in segmented point cloud

Many studies and solutions have recently been looked into to recognise and classify objects from 3D point clouds [Aldoma et al., 2012, Fernandes et al., 2021, Rutzinger et al., 2008]. Many open sources datasets have been created and published like the *Mian* dataset [Mian et al., 2006] (figure 4.4) or *ShapeNet* dataset [Chang et al., 2015] for benchmarking algorithms. Among the main approaches to classify objects from point clouds, the main ones are:

- (i) *semantic point cloud descriptors* which are inspired by the human perception of an object: by touching an object and without seeing it, a human is able to guess the nature of the object. The touch only reveals the object's shape and curvature,

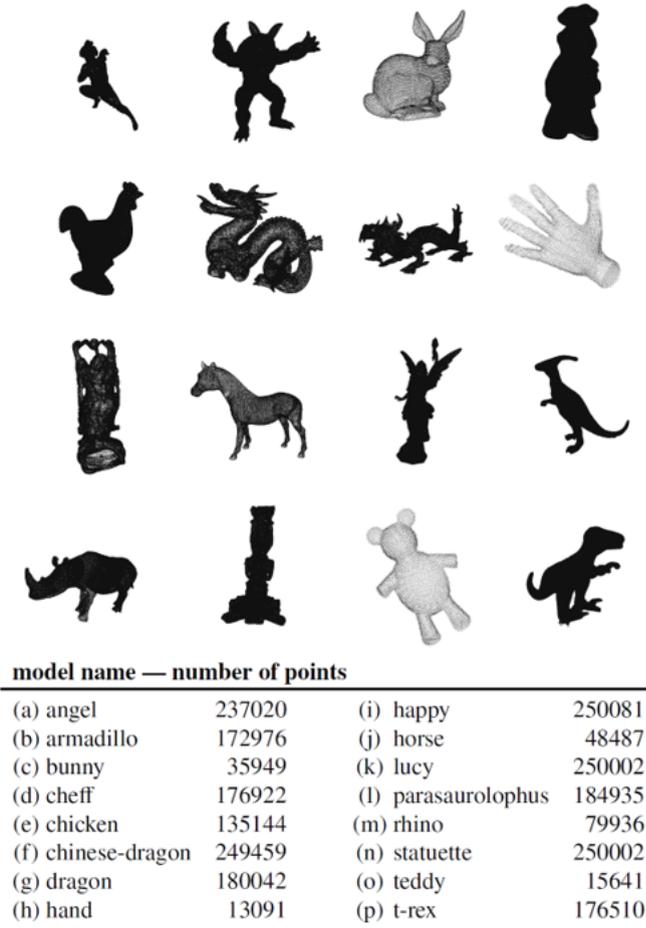


FIGURE 4.4: Reference point cloud objects and labels from the *Mian dataset*. From Lev [2020].

hence the idea behind this algorithm is to use a measure of local curvature to classify an object [Lev, 2020].

- (ii) *deep learning networks* such as PointNet [Qi et al., 2017] that directly consumes point clouds to produce object classification and semantic parsing with good accuracy but requires a labelled dataset with objects associated with masks and heavy computation.

Seismic point cloud objects pose specific problems. First, we do not have a labelled dataset necessary to train a neural network. Secondly, the seismic objects are all relatively flat compared to everyday life objects (figure 4.4). It is insufficient to characterise seismic objects based solely on curvature descriptors. Finally, the diversity of shapes of turbidite fans (figure 4.5) makes it complex to define logical, hard-coded rules for classification.

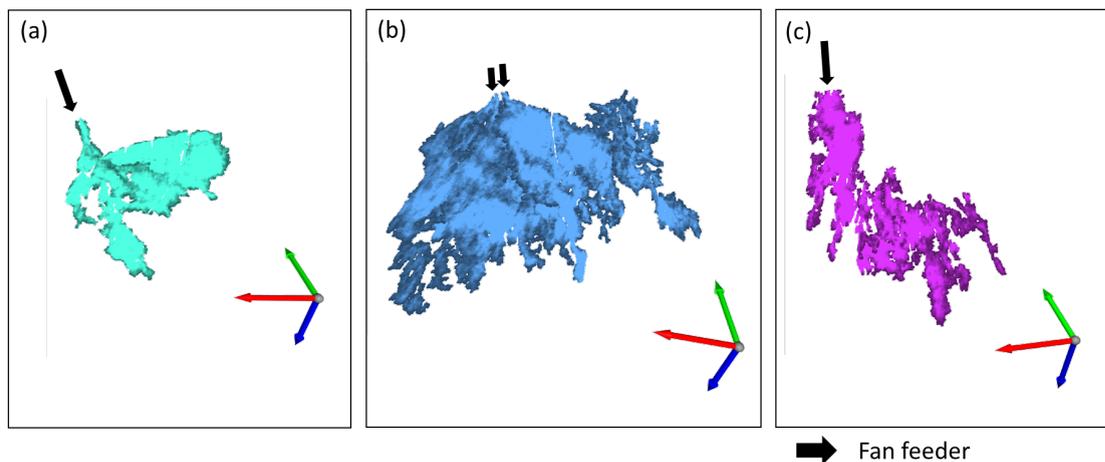


FIGURE 4.5: Illustration of the diversity of shapes and geometries between turbidite fans. Even amongst the three fans observed in our *validation dataset*, the fans vary significantly in terms of length, stretch or aspect of input feeders.

For instance, looking only at the fan feeders of the three fans originally delineated in Figure 4.5, we can see quite a diversity: Figure 4.5.a, the feeder is well preserved, and we can observe the channel coming down the slope, Figure 4.5.b, the feeder is less preserved, and we observe a double entry feeder, and Figure 4.5.c, the geometry of the feeder is not preserved, and we can only guess it from the slope orientation and the amplitude brightness.

4.2.2 Proposed workflow for the detection of geobodies among point cloud objects

We propose a two-step workflow for detecting geobodies from point cloud objects (figure 4.6). The first step is feature engineering to extract the object characteristics (e.g. amplitude mean, aspect ratio), and the second step retrieves particular geobodies based on feature similarity.

The feature engineering allows the user to select, control and understand the descriptors used to characterise the objects. Thus, it provides some flexibility for the users to adapt the feature extraction to the classification task or type of geobodies they want.

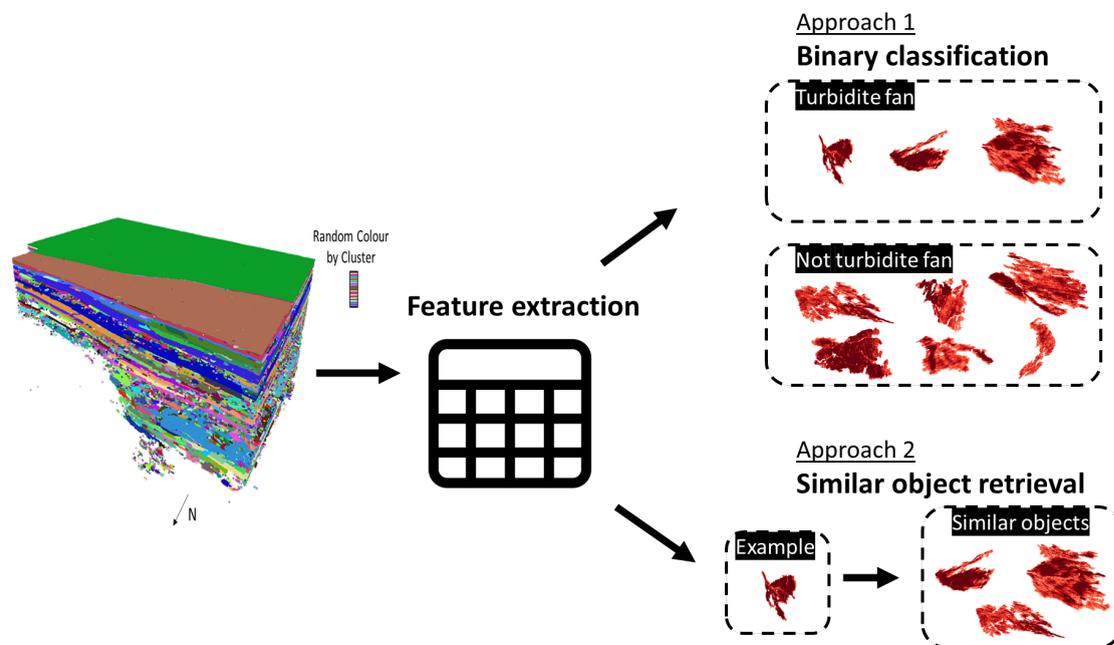


FIGURE 4.6: Object detection workflow from a segmented seismic.

For the second step, we investigate geobody detection by a closest object retrieval approach: given an input example, the solution retrieves the most similar objects to this example.

In this chapter, the feature engineering and extraction are presented first, and then, the geobody retrieval is introduced and assessed using our *validation dataset*.

4.3 Feature engineering and data analysis

As mentioned in chapter 1 Introduction, a seismic interpreter confidently detects turbidite fans through two types of criteria: (i) the particular shape of the fans: a thin feeder coming down a slope and larger lobes extending from the feeder, and (ii) the amplitude distribution on the geobody with bright amplitudes going down the feeder to the centre of the lobes and weaker amplitudes on the edges and lobe terminations. These criteria guide us in the selection of the features to extract.

Feature extraction refers to transforming raw data into numerical descriptors (features). Features are pertinent information that best describes the data. Working with features

produces better and more interpretable results than directly applying machine learning to the raw data. For each object or cluster in the segmented seismic, we perform an extraction of the same features, transforming the point cloud object into a list of feature values. At the end of the feature extraction process, we obtain a table where each row represents a cluster, and each column is a feature.

This section presents first an overview of the building of the features (4.2.1 & 2), a short summary of these features and the preprocessing we apply (4.2.3), and data analysis (4.2.4).

4.3.1 Geometric shape features

We extract from the point cloud five types of geometric descriptors: size, contour ratio, aspect ratio, elevation distribution and geometrical seismic attributes.

Size

As a proxy to the object's size, we extract the number of points in the point cloud object ($N(points)$).

$$N(points) \tag{4.1}$$

Another possible size feature is the object's surface on an XY plane (S). Compared to $N(points)$, it eliminates the overlapped points in the point cloud.

Contour ratio

The contour ratio is the ratio of the surface of the object over the length of its contours in 2D ($r_{contours}$), contour meaning the edge of an object. The surface of the object and the length of the contour are expressed in the number of pixels. To extract the contour of the point cloud object, we use a standard image processing procedure consisting of 3 steps (figure 4.7). First, we perform a vertical projection of the point cloud in 2D, creating an image. We only record at each pixel of the image the presence or absence of points (0 or 1). Then, we blur the image to eliminate small rough edges and holes.

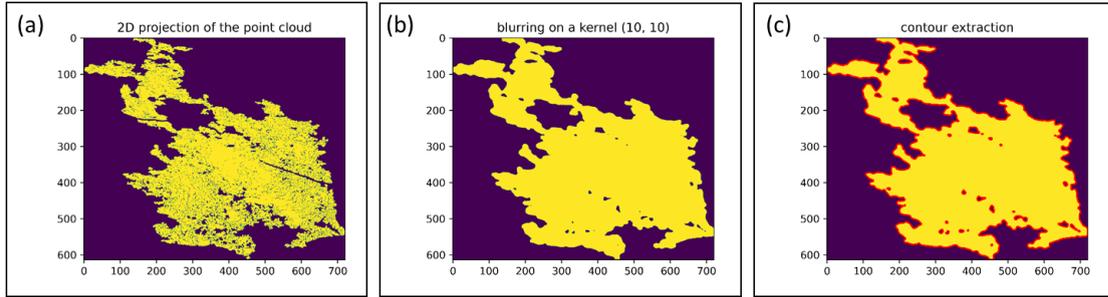


FIGURE 4.7: Workflow to extract the contours of an object in a 2D projection. (a) step 1: 2D projection; (b) step2: blurring to eliminate rough edges and small holes; (c) step 3: contour extraction.

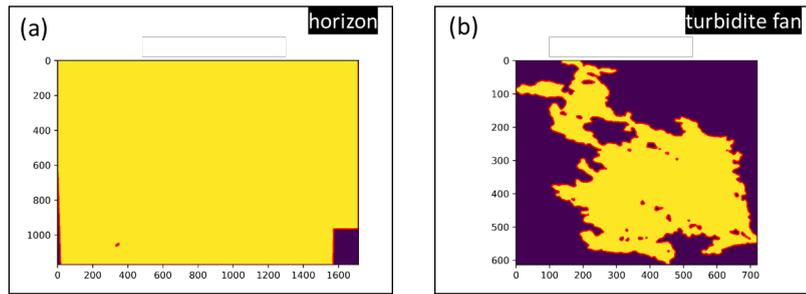


FIGURE 4.8: Comparison of $r_{contours}$ between a horizon (a) and a turbidite fan (b).

The Blurring is done by convolving an image with a normalized box filter. It simply takes the average of all the pixels under the kernel area, rounds it to the closest value (0 or 1) and replaces the central element. We use a (10,10) kernel in this work. Finally, we extract the contours from the binary image using the topological structural analysis of digitized binary images by border following algorithm [Bradski and Kaehler, 2008, Suzuki and be, 1985].

To get $r_{contours}$, we then simply divide the length of the contours by the surface, all expressed in number of pixels.

$$r_{contours} = \frac{\sum d_{contour}}{S} \quad (4.2)$$

Intuitively, $r_{contours}$ of a turbidite fan will be much higher than the one of a horizon due to the complexity and roughness of its edges and the smaller area of the object (figure 4.8).

Aspect ratio - Eigen-based features

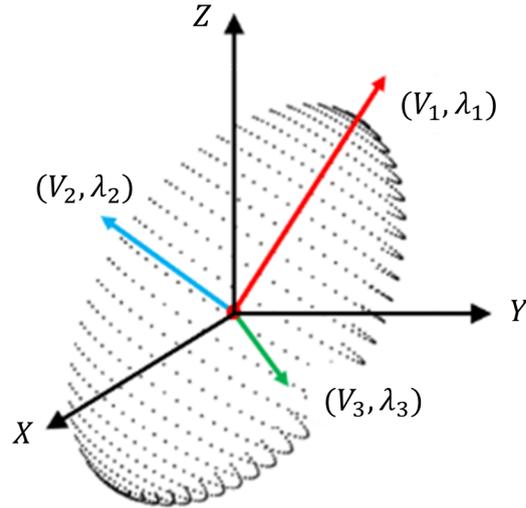


FIGURE 4.9: Diagram of point cloud eigenvectors and eigenvalues. λ_i is the i^{th} eigenvalue and V_i is the i^{th} eigenvector Figure from Chen and Lin [2019].

Features based on the aspect ratio in 3D of the object are common features used to describe a point cloud [Bassier et al., 2020, Grilli et al., 2021]. They are derived from the point cloud's eigenvalues (figure 4.9). The most common eigen-based features include (table 4.1):

Feature	Equation
Linearity	$\frac{\lambda_1 - \lambda_2}{\lambda_1}$
Planarity	$\frac{\lambda_2 - \lambda_3}{\lambda_1}$
Sphericity	$\frac{\lambda_3}{\lambda_1}$
Slope	$\arccos\left(\frac{V_{1,3}}{\sqrt{V_{1,1}^2 + V_{1,2}^2 + V_{1,3}^2}}\right)$
Omnivariance	$\sqrt[3]{\prod_{j=1}^3 \lambda_j}$
Anisotropy	$\frac{\lambda_1 - \lambda_3}{\lambda_1}$
Eigenentropy	$-\sum_{j=1}^3 \lambda_j \ln(\lambda_j)$
Sum of eigenvalues	$\sum_{j=1}^3 \lambda_j$
Surface variation	$\frac{\lambda_3}{\sum \lambda}$

TABLE 4.1: Eigen-based features.

To describe our seismic objects, we kept only *linearity*, *planarity* and *sphericity*.

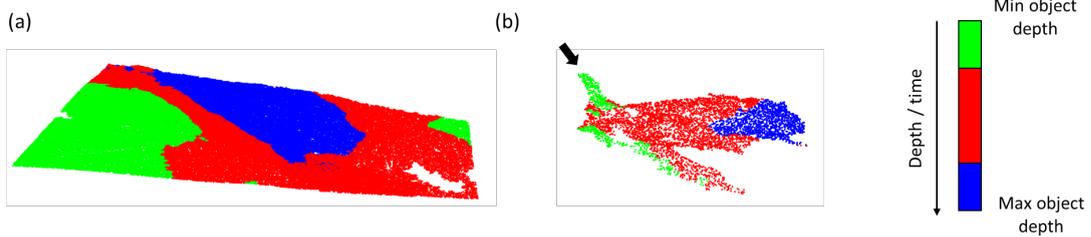


FIGURE 4.10: Illustration of the depth distribution of a horizon (a) and a fan (b). Points are coloured according to two depth boundaries, above the first elevation quarter and below the last quarter. The ratio of points of each elevation bin is measured.

The main *orientation* of the object can also be derived from the eigenvectors of the point cloud (orientation angle derived from the main eigenvector – equation 4.3).

$$orientation = atan2(V_{1,y}, V_{1,x}) \quad (4.3)$$

orientation can be a pertinent feature in some cases, for instance, when looking for turbidite fans from a particular slope with a specific orientation.

Elevation distribution

The depth distribution is characteristic of geological objects. While most geological structures are planar and predominantly horizontal, fans and sills are only partially so. Fans are distinguished by having a feeder above the lobes. Similarly, sills are not totally planar; their terminations tend to rise vertically, giving them a saucer-shaped appearance. Thus, we measure the ratio of points to two depth boundaries, above the first quarter and below the last quarter (figure 4.10), and store it as two additional features.

Geometrical seismic attributes: mean dip, mean azimuth and mean curvature

Common geometrical attributes like dip, azimuth or curvature characterise the local structure of the features in the seismic. The mean value of each of these attributes on the point cloud object provides information on its structure.

Other interesting geometrical features would be features derived from the comparison between neighbouring point clouds. For instance, the mean local dip of an object could be compared to the mean local dip of surrounding objects. It could highlight discordant

objects and unconformities in the seismic data and thus better distinguish discordant geobodies such as igneous intrusions.

4.3.2 Amplitude-based features

We extract from the point cloud four types of amplitude-derived descriptors: mean amplitude, mean coherence, mean Zeboudj's distance and other statistical measures of textural characteristics.

Mean amplitude

The mean of the seismic amplitude of the object.

Mean coherence

The mean of the seismic coherence of the object. In this study, we calculated the coherence mean as coherence was already computed for the point cloud extraction (chapter 3 – 3.2.2 Creation of a seismic data point cloud), but other seismic attributes can be considered. For instance, frequency attributes would help characterise specific geological patterns that are not observable directly from the amplitude attribute but become apparent when the signal is decomposed at certain frequency bands.

Mean Zeboudj's distance

The Zeboudj's internal disparity (CI), presented in chapter 3 – Equation 3.7, characterises the spatial distribution of a property by distinguishing between e.g. homogeneous and heterogeneous amplitude distributions (or any other seismic attribute on the point cloud). Furthermore, the sphere radius δ can be varied to create different features characterising the differences of any property over different distances.

Other statistical measures of textural characteristics

Texture analysis aims to mathematically describe how pixel values are distributed across an image or point cloud. Zeboudj's distance is one approach, but other statistical measures can characterise the spatial distribution of a property (e.g. amplitude).

Widely used in geostatistics, the range of the semi-variogram indicates the distance from which two points (two-point statistics) are no longer spatially correlated - the distance after which the variogram model levels off [Matheron, 1963]. However, computing the variogram has a higher computational cost, $o(n^2)$ complexity.

Another texture analysis technique extensively used in image processing is the grey-level co-occurrence matrix (GLCM). First introduced by Haralick et al. [1973], the GLCM computes the distribution of co-occurring pixel values at a given offset. Haralick et al. [1973] proposed fourteen different measures of textural features derivated from GLCMs, but only five are commonly used and suffice to give good results in a classification task [Patil et al., 2011] (table 4.2). Each of these features represents a specific property of the texture. The Haralick features have already been successfully applied to classify seismic textures [Chopra and Alexeev, 2006, Gao, 2003, Le Bouteiller and Charléty, 2020, Le Bouteiller et al., 2019], and proven efficient for seismic interpretation. However, computing GLCMs is computationally demanding and highly sensitive to the number of grey levels (amplitude bins), $o(n^2)$ [Clausi and Jernigan, 1998]. But a more suitable computational strategy proposed by Clausi and Zhao [2003] dramatically reduces the completion times, making it more suitable for operational use.

Haralick Features	Equation
Homogeneity	$E = \sum_i \sum_j (M(i, j))^2$
Contrast	$C = \sum_{k=0}^{m-1} k^2 \sum_{ i-j =k} (M(i, j))$
Correlation	$Cor = \sum_i \sum_j \frac{(i-\mu_i)(j-\mu_j)M(i,j)}{\sigma_i\sigma_j}$
Entropy	$H = \sum_i \sum_j M(i, j) \log(M(i, j))$
Local homogeneity	$LH = \sum_i \sum_j \frac{M(i,j)}{1+(i-j)^2}$

TABLE 4.2: List of the five most common Haralick features, M being the GLCM. Each of these features represents a specific property of the texture of a property (e.g. amplitude).

4.3.3 Feature selection and preprocessing

The feature extraction process results in a total of 24 features describing different properties of the objects in terms of their geometries and amplitude textures. This list of features is not exhaustive but represents commonly used descriptors for 3D objects and texture analysis. Not all 24 features should be used for any interpretation task, but the features should be selected to fit the characteristics of the desired geobody. The selection of features can be made by observing the distribution of the features and the feature's ability to isolate the desired geobody, but the selection should also be made by using common sense on whether the feature corresponds to a property that could potentially help the classification of the geobody. For instance, for fans, a small slope is an important characteristic of distinguishing fans. The advantages of using a feature extraction pipeline compared to a deep neural network are that the features have a meaning, and it provides a control to the interpreter to select which characteristics to use.

Before studying the feature space of a dataset, it is common to apply some preprocessing steps to transform raw feature vectors into a representation more appropriate for any downstream estimator -classifier, clustering algorithm or similar- [Pedregosa et al., 2011]. For our application, we apply standardisation to the features by removing the mean and scaling the features to unit variance, which makes the representation more appropriate and balanced between features. However, other scaling methods could be considered, such as normalisation, min-max scaler or robust scaler [Pedregosa et al., 2011].

4.3.4 Data analysis

The objective of the data analysis is to understand how the extracted features are distributed, analyse the distribution of the geobodies amongst the other objects and potentially find a combination of features that bring together and isolate the desired geobodies.

We carried out this analysis on the *validation dataset* and took as fans the three fans manually interpreted at the initial stage (figure 4.5). We filtered the dataset to contain

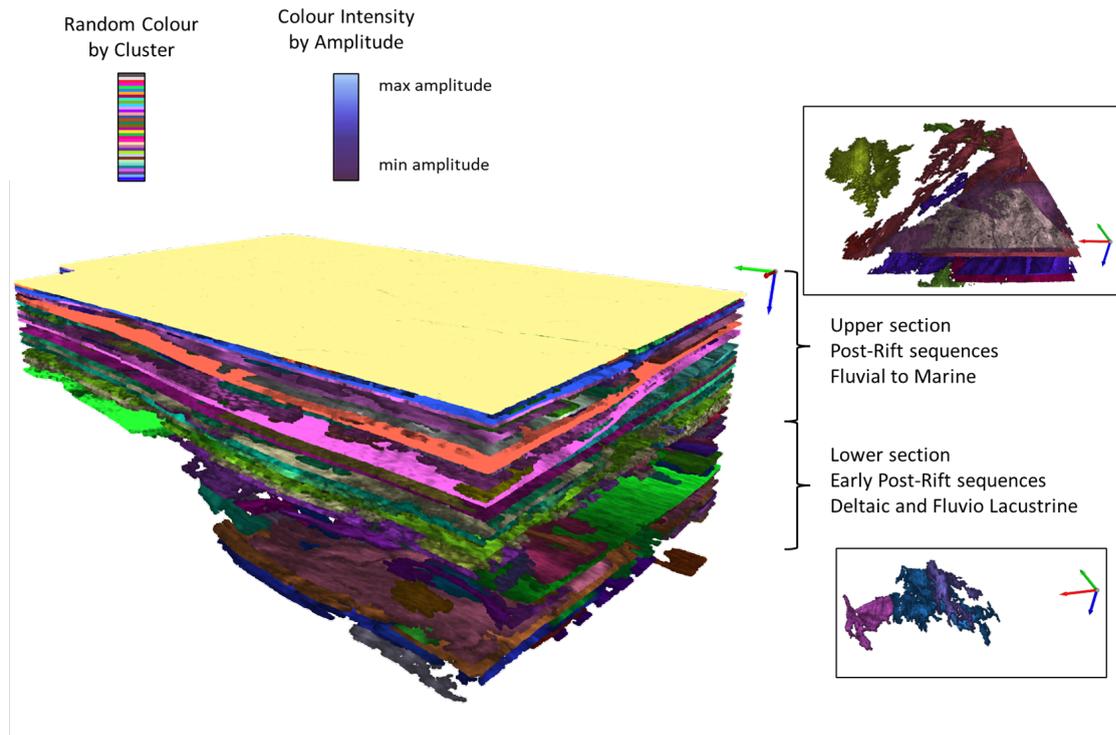


FIGURE 4.11: Visualisation of the main objects present in the *validation dataset*. The upper section of the dataset contains mainly extended horizons and fluvial systems. The lower section contains the infills of the lake with some deltaic deposits and fans. To enhance the visualisation of the segmented seismic, we use the HSV (Hue, Saturation, Value) colour representation where hue and saturation represent the random colour of the cluster, and the value represents the seismic reflection amplitude.

only clusters with more than 5000 points. It removed the clusters that were too small to represent any geobody of interest. Thus, the dataset resulted in 394 clusters (objects) representing a good diversity of potential objects present in the seismic data. The objects can be divided into two main packages: (i) the upper section, which mainly contains extended horizons and fluvial and channelized systems, and (ii) the lower section, which contains the infills of the lake with some deltaic deposits and fans (figure 4.11). Three fans have been manually interpreted and were used to analyse the relevance of the features, but additional fans, present in the lower section of the *validation dataset*, were not labelled.

To identify fans, we extracted seven features: the mean amplitude, the mean coherence, the mean Zeboudj's distance, the number of points, the contour ratio, the linearity and the sphericity for each of these 394 objects. The choice of the features is made considering (i) the most intuitive descriptors for turbidite fans, (ii) keeping balanced numbers between

shape-based and amplitude-based features, and (iii) keeping a limited number of features to facilitate the analysis. The purpose here is to illustrate some characteristics of fans, but these choices should be considered differently for every interpretation task.

4.3.4.1 Univariate and bivariate analysis

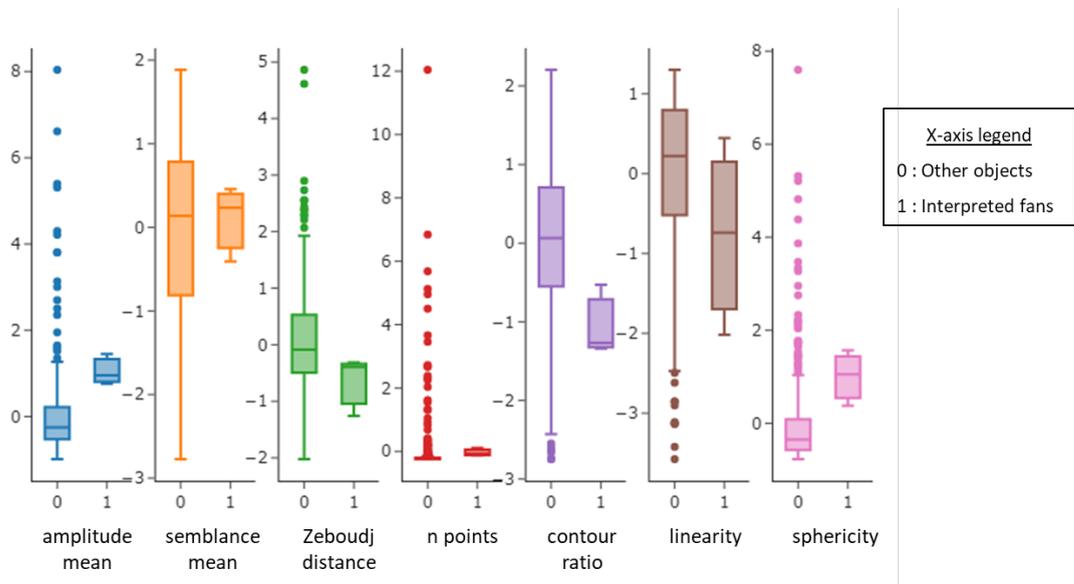


FIGURE 4.12: Box-plot of the selected normalized features for fan retrieval. The distribution of the three interpreted fans is compared to the other objects delineated in the *validation dataset*.

The box plot analysis gives a first idea of the object distribution for all features. In Figure 4.12, we compare the distribution of each feature for the three interpreted fans (i.e. 3 data points) vs the rest of the data set. We observe that the distribution of the fans is restricted for most features but is wider for linearity, suggesting that fan linearity can vary greatly. Looking at *npoints*, most objects have a relatively low number of points, and the three fans have a relatively high number of points. Also, the points are not optimally spread through the values of each feature. Other scaling preprocessing solutions more appropriate to deal with outliers should produce better results.

Cross-plot matrices help to better understand the object distribution across the feature space and identify potential relationships between features. It provides a quick representation of each feature and bivariate analyses by pair of features (figure 4.13). Having only three fans labelled among the 394 objects and potentially several non-labelled fans

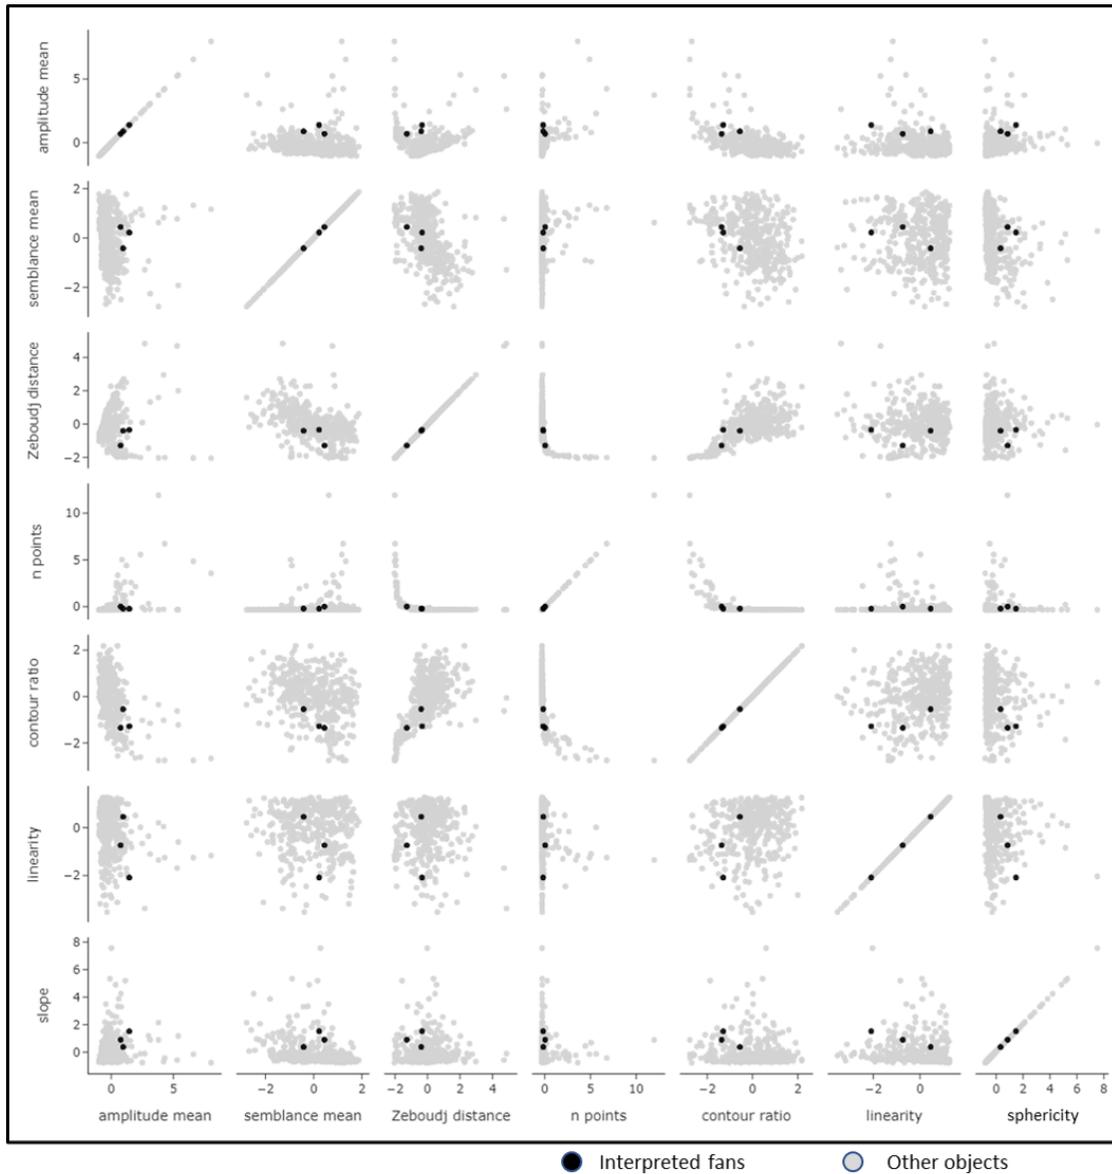


FIGURE 4.13: Cross-plot matrix of the selected features for fan retrieval. Each point represents an object in the segmented seismic. The three fans manually interpreted are plotted in yellow, and the other objects in brown.

makes it difficult to draw any conclusion. Still, by looking more precisely at the cross-plots of the mean amplitude versus the contour ratio and of the sphericity versus the contour ratio (figure 4.14), we can notice that the three fans are grouped together and separated from most of the other objects. This indicates that these attributes are likely to be good indicators of turbidite fans.

Fans are slope deposition elements, so it seems relevant to focus on the slope of each object or the sphericity. Here we focused on the sphericity value of the objects. We filter

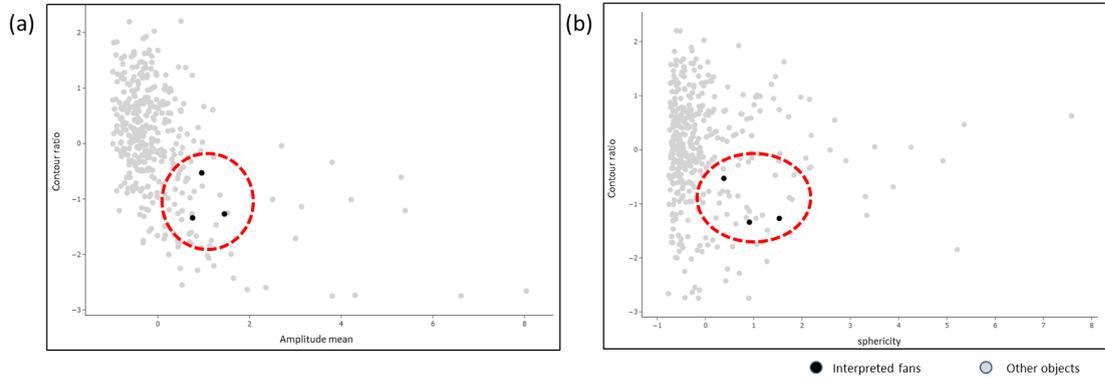


FIGURE 4.14: In particular, the cross-plots of contour ratio vs amplitude mean and contour ratio vs sphericity allow us to distinguish and group the three manually interpreted fans. This result suggests that the amplitude mean, the contour ratio and the sphericity of the objects are pertinent features to characterise fans.

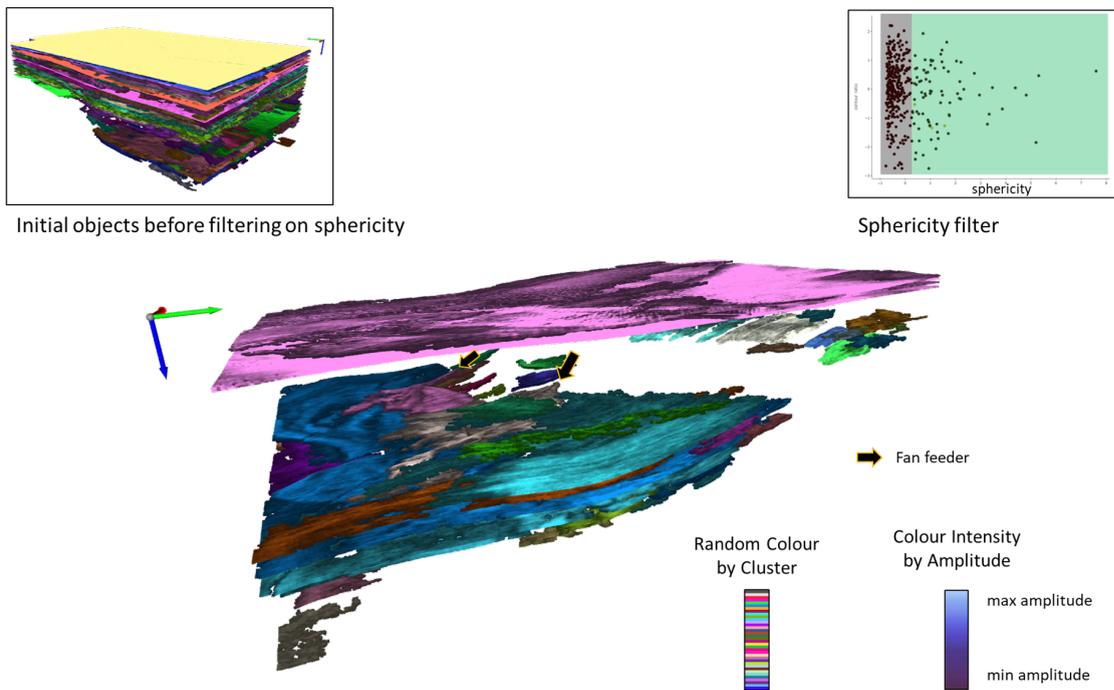


FIGURE 4.15: 3D visualisation of objects with high sphericity, displaying steep slopes or overlapped structures. The steep objects are mainly found in the lower section of the dataset, where structures are constrained and deposited along the lake slope.

by the objects by sphericity and plot only the objects with high sphericity (figure 4.15). From the analysis of steep objects in the dataset, we observe (i) most upper section objects where no fans can be deposited are filtered, (ii) two of the three fans are still present, but the third one is not – its lower sphericity value is explained by the fact that the feeder of this third fan is less preserved, and (iii) one horizon from the upper section is still present. The presence of the horizon in the upper section is due to the overlapping structure in the horizon, impacting the vertical component of the eigenvalues (λ_3) in the sphericity value. From these observations, sphericity seems a relevant feature to help characterise fans, but poorly preserved or overlapping structures add complexity to that analysis. One thing to consider is that fans are not the only objects that might have a slope or high sphericity. For instance, if dykes are also present in the dataset, they will have strong slope values. Thus, having a prior knowledge of the characteristics of the geobody to look for and of what other types of geobody might be present is essential, and geobody detection should not be based on a single feature but on the combination of multiple features.

4.3.4.2 Multivariate analysis - dimension reduction

To further analyse the 7-dimensional space, we used a dimensional reduction technique called t-SNE (*t-distributed stochastic neighbour embedding*) to visualise the distribution of objects in the feature space in 2D (figure 4.16).

T-SNE is a nonlinear dimensionality reduction technique for visualising high-dimensional data projected on a low-dimensional space and where the trends of the original dimensional space are preserved [Hinton and Roweis, 2002, Van Der Maaten and Hinton, 2008]. First, t-SNE constructs joint probability distributions over pairs of objects so that the probabilities reflect the similarities between the data points. Then, t-SNE attempts to minimise the Kullback-Leibler divergence between the joint probabilities of the high-dimensional data and the low-dimensional embedding [Pedregosa et al., 2011].

The analysis of the t-SNE plot of the seven features suggests that the extracted features are relevant to characterise fans with the relative proximity of the three interpreted fans in the low-dimensional embedding (figure 4.16). The perplexity value for the t-SNE

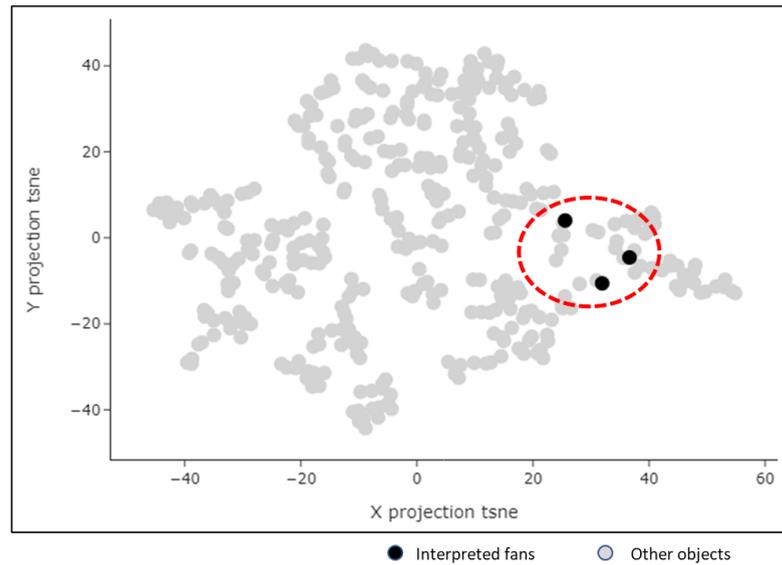


FIGURE 4.16: t-SNE representation of the seven extracted features in a 2D projection. It shows proximity between the three manually interpreted fans which implies similarity and proximity of the fans in the 7-dimensional feature space.

plot used in figure 4.16 is ten which is included in the recommended range of perplexity values (5–50) suggested by Van Der Maaten and Hinton.

4.3.4.3 Summary of observations from the data analysis

The main observations from the data analysis are:

- (i) The seven features extracted - mean amplitude, mean coherence, mean Zeboudj distance, number of points, contour ratio, linearity and sphericity - seem appropriate to characterise the general properties of turbidite fans. However, the feature extraction could still be improved, and the selection should be reconsidered for any different type of detection.
- (ii) It is hard to define clear, robust classification rules based on the characteristics of the fans. For instance, high sphericity of the object tends to highlight fans, but this is only when the feeder is preserved.
- (iii) Having only a few labelled fan examples and several potential non-labelled fans, caution should be exercised before drawing any general conclusions, but further tests will be carried out in Chapter 5.

4.4 Geobody detection by closest object retrieval

In this section, we first introduce the methodology used for the geobody detection, then we present the results obtained on the *validation dataset*, and finally, we assess the methodology and discuss the advantages and potential risks.

4.4.1 Methodology

Before presenting the final methodology used for our geobody detection method, we first present the initial attempts at geobody detection that we tested and their limitations as it justifies our choice for our final methodology: geobody detection by closest object retrieval.

4.4.1.1 Initial attempts at the geobody detection

The first intuitive ML approach to detect a geobody is to frame the problem as a binary classification where an algorithm learns in a supervised way to classify each object as either **Turbidite** (1) or **Not Turbidite** (0). With the features of each object extracted, we now consider a tabular data structure. Supervised machine learning offers several algorithms to solve classification problems from tabular data, the most popular being Support Vector Machine (SVM), K-Nearest Neighbour (KNN), Decision Tree and other tree-based models (Random Forest, XGBoost...). However, in order to implement this, the fan classification poses a number of problems [Shwartz-Ziv and Armon, 2022] common in seismic interpretation.

The three main problems for the geobody classification are:

- (i) The low number of training samples(394 objects), which is very small to train a model in a supervised way. A limited number of training samples is a common problem in seismic interpretation tasks: although more data getting publicly available, finding labelled dataset remains problematic.

A standard solution is to artificially create more training samples with some data augmentation procedure - rotating the object, resizing the object, and adding noise points to the objects -. However, these data augmentation operations would not increase the dataset regarding the diversity of fans -different linearity, shapes, types of feeders, etc.-.

- (ii) the highly imbalanced classes. The second issue is the highly imbalanced number of samples per class between Turbidite (3 objects) and Not Turbidite (391 objects).

Two main strategies can be used to cope with imbalanced classes: (i) under-sampling, which consists of deleting instances from the over-represented class, but it is not recommended in this case with already too few samples, and (ii) over-sampling, which increases the number of rare samples. Over-sampling synthesizes new examples close to the original minority class samples. The most common over-sampling technique is Synthetic Minority Oversampling Technique (SMOTE) [Chawla et al., 2002]. Otherwise, to deal with unbalanced classes, the chosen machine learning models need to incorporate a class weighting in their loss function to ensure that all classes are represented. Decision trees often perform well on unbalanced datasets, as the splitting rules used to create the trees take into account the class variable and force both classes to be addressed.

- (iii) the false-negative labels in the training dataset. Finally, the third main issue is the fans present in our dataset but not labelled as fans. Indeed, only three fans were manually interpreted, but more fans are present in this cropped seismic. The number of fans present is unknown but is likely to be of a high proportion compared to the only three labelled fans, which will negatively impact the training. One easy solution would be to interpret and label the remaining fans manually. However, false negative labels are common in seismic interpretation, and we want a general solution to this problem.

The traditional classification ML approach was the first path explored to resolve our object detection problem. However, that approach was not fitting the constraints of our problems and data. The lack of a proper training dataset forced us to change our approach to this problem.

We decided to change how we posed the problem: instead of seeing it as a classification problem, we chose to see it as a closest-object-retrieval problem. This new approach is presented in the following section and used in the rest of the thesis.

4.4.1.2 Closest object retrieval paradigm

Principle

The closest object retrieval approach finds a set of n nearest-neighbours: objects within a seismic volume that are most similar to an object of interest provided by the user. The interpreter may choose the example geobody in different ways: (i) it may be a geological body of particular interest that the interpreter identified by first scoping through the segmented seismic data, (ii) it may be a geological body previously identified in another seismic dataset and input for retrieval in newly segmented seismic data, or (iii) it may even be a synthetic object that has the idealised shape the interpreter wants to find.

Advantages

The main advantages of using a closet-object-retrieval are:

- (i) No training dataset is required, solving our lack of data issues. This hybrid approach does not require supervised training of the models. Only the position of the objects in the feature space is explored, looking for similarities to the object example.
- (ii) It provides more precision on the characteristics to look for in feature space by considering the additional information provided by the object example.
- (iii) More control is provided to the user. By inputting a specific geobody as the example to look for in the dataset, the user has a direct lever to direct the search for geobodies.
- (iv) It offers greater adaptability to a change in the dataset. As the approach is unsupervised, the detection is not trained on a specific seismic dataset. Therefore, there is no introduction of a bias towards specific processing or resolution of the seismic data. For each dataset, the features of the objects are explored from their own

aspects, and the input example can be introduced knowing the specificity of the explored seismic data.

- (v) Finally, it offers better adaptability to different types of objects without requiring new training for this new classification. Again the only supervision to the detection being the example input, a change of example will mean a change in the type of object searched. What is critical to adapt to the type of objects searched for is the feature selection, so the feature space accordingly represents the important object characteristics.

The main problem with similarity searching is efficiency and scaling with high-dimensional attribute spaces. For a representation of objects as points in high-dimensional attribute space, a similarity search is a nearest-neighbour query problem. A brute-force linear search approach would solve the problem but would not scale well with a high number of points or high dimensional spaces. A k-d tree structure would do better in terms of performance, but would still involve the inspection of a large proportion of the database at each query. Gionis et al. proposed in 1999 Locality Sensitivity Hashing (LSH) [Indyk and Motwani, 1998] to approximate nearest-neighbour search in an efficient manner. Thus, in this work, we use LSH to perform geobody retrieval in an efficient manner.

4.4.1.3 Binary hashing for fast object retrieval with LSH

LSH or bucket search consists in hashing the points in the database and associating them to a binary vector representation with a high probability that similar binary codes are mapped to similar points [Gionis et al., 1999]. A number of hashing algorithms have been proposed to approximate similar objects to an input query [Gionis et al., 1999, Gong et al., 2013, Weiss et al., 2008]. LSH hashes similar points in the feature space into buckets. A bucket is a subdivision of the space and is identified by a binary code. During the query phase, only the points in the same bucket are inspected, and the nearest ones are returned, thus avoiding an inspection of the entire database. Many different implementations for LSH exist, but we are using locality sensitive hashing with persistence support. Figure 4.17 illustrate an LSH with persistence support implementation on a simple example.

Each object is represented as a point in the feature space. First, a number of hyperplanes are created through the feature space. Then, a binary code is constructed for each point of the database based on being above or below each hyperplane; each digit of the binary code represents the point position to a specific hyperplane (0:below, 1:above). The space is thus divided into sections (Equiv. buckets) identified by a unique binary code. At the query stage, a binary code is constructed for the new point (example provided by the user) following the same principle. The new query object is then only compared to other objects with the same binary code for similarity distance in the feature space. In practice, several hashing tables with binary codes are created to ensure a more accurate approximation.

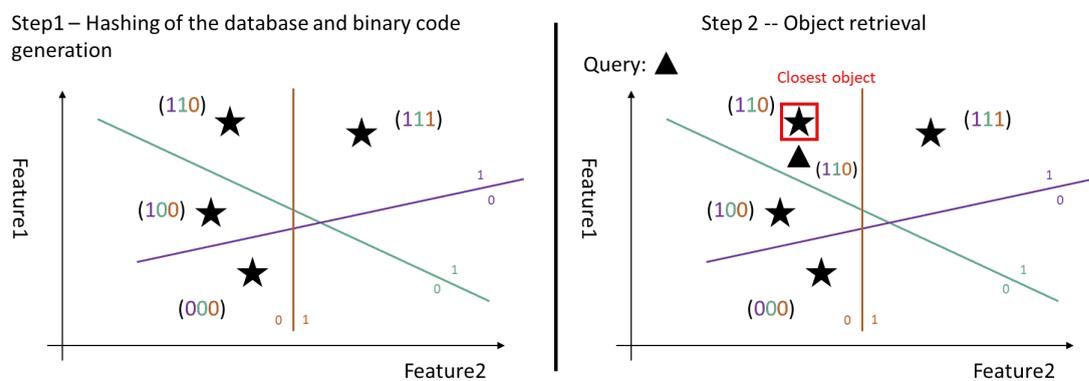


FIGURE 4.17: Illustration of locality hashing of the attribute space and retrieval of the closest object. We illustrate LSH on a two-dimensional attribute space example with a database of four objects (stars). In the first step, three hyperplanes are randomly created and divide the attribute space. The position of each point of the database is evaluated as below (0) or above (1) the hyperplane. Thus, a binary sequence of the approximate position of each point gets created. In the second step, a query object is entered (triangle), looking for similar objects. The new point gets hashed using the same methods, and the object with the same binary code is returned.

Recently, LSH has been combined with deep neural networks and applied to successfully retrieve similar images in a database [Lin et al., 2015, Liu et al., 2016, Zhu et al., 2016]. They used convolutional neural networks to extract the hash representations of the images and demonstrated scalability and efficacy on a large-scale dataset of millions of images.

4.4.2 Results

To first validate the method, we apply it to the *validation dataset* for fan retrieval. The methodology consists first in hashing the database of 394 objects to binary codes. Then, one of the three manually interpreted fans is shown as an example of similar objects, asking for the 30 closest objects. The first validation step is to see if the method was able to retrieve the two other interpreted fans, and then the second validation step is to see if other potential fans were discovered.

As the number of objects studied and the dimensionality of the attribute space are relatively small, the hashing and querying of the closest objects take place in near-to-real time.

The results of the experiment are shown in Figure 4.18. The first observation is that the extended horizons are filtered out, making observing subtle objects easier (figure 4.18(c)). Secondly, the three interpreted fans are successfully retrieved amongst the 30 objects. When looking more precisely at the location of the fans (figure 4.18(d)), other objects showing depositional patterns and fan-like shapes are observed amongst the 30 filtered. In particular, one object is channelised (figure 4.18(d) - object D), and a feeder can be distinguished. Overall, while many of these objects are clearly not fans, several have shapes and patterns that might suggest potential fans and merit further study. Even the few objects retrieved from the upper section present channelised patterns with complex shapes.

4.4.3 Assessment of the methodology

These initial results highlight the potential of the closest object retrieval approach as a means of rapidly interrogating a database of seismic objects and extracting interesting targets. Our workflow still requires the interpreter's expertise to sort and validate relevant objects from the returned ensemble, but we benefit from a reduced number of clearly delineated objects to evaluate.

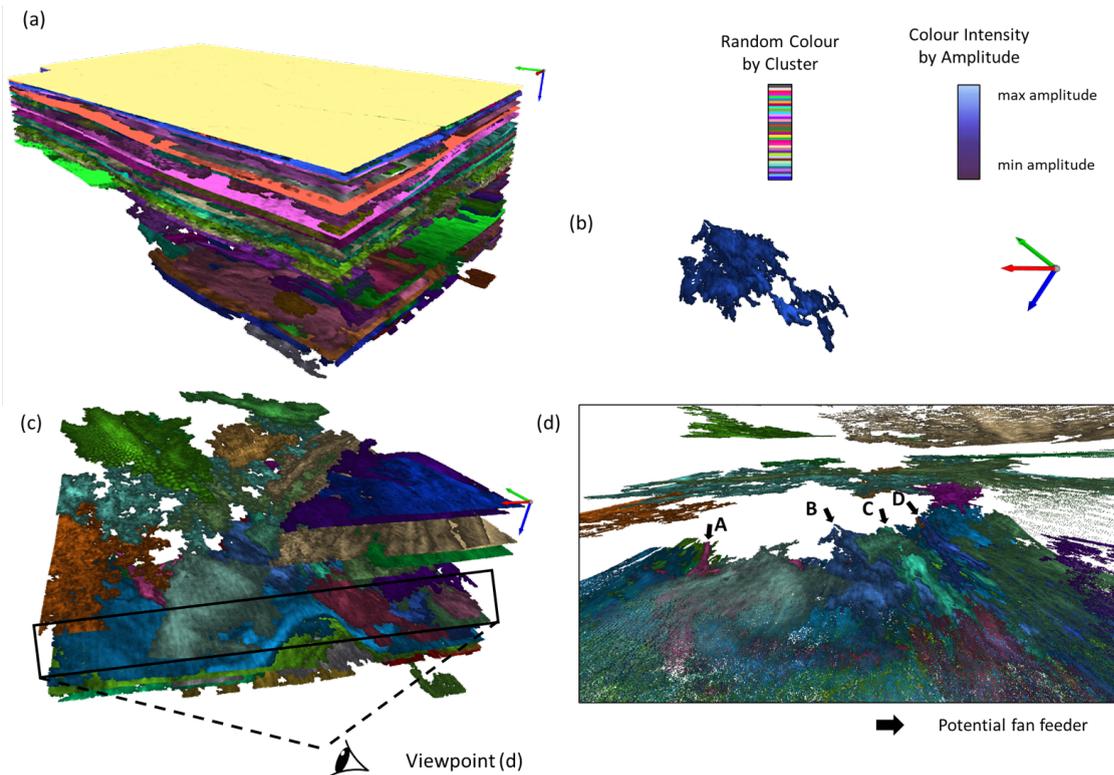


FIGURE 4.18: 3D view of the results of the 30 most similar objects to a fan on the *validation dataset*. The different elements are (a) The database of 394 seismic objects; (b) the fan example shown for retrieval; (c) the results of the 30 most similar objects; and (d) a zoom-in of (c) to the location of the main fans, and where we can observe the fan feeders going down the slope.

A risk exists that a potential geobody will be dropped from the subset of retrieved objects if the geobody is too far away in the feature space from the example object. This risk underlines the importance of the choice of the feature extracted to accurately characterise the type of geobody.

The results depend strongly on the object inserted as an example and its characteristics. We vary the input fan to test the effects and assess the robustness of the fan retrieval. Figure 4.19 compares the results for each of the three manually interpreted fans. Depending on the fan, different objects get output. With example three (figure 4.19(d)), the three interpreted fans are present, but not with examples 1 and 2 (figure 4.19(b) and (c)). The results in example 3 (figure 4.19(d)) illustrate the connection between the characteristics of the example and the objects retrieved: the objects all have complex elongated shapes similar to the example fan (figure 4.20). In addition, in Figure 4.19(d), two potential fans are observed with elongated and channelised shapes.

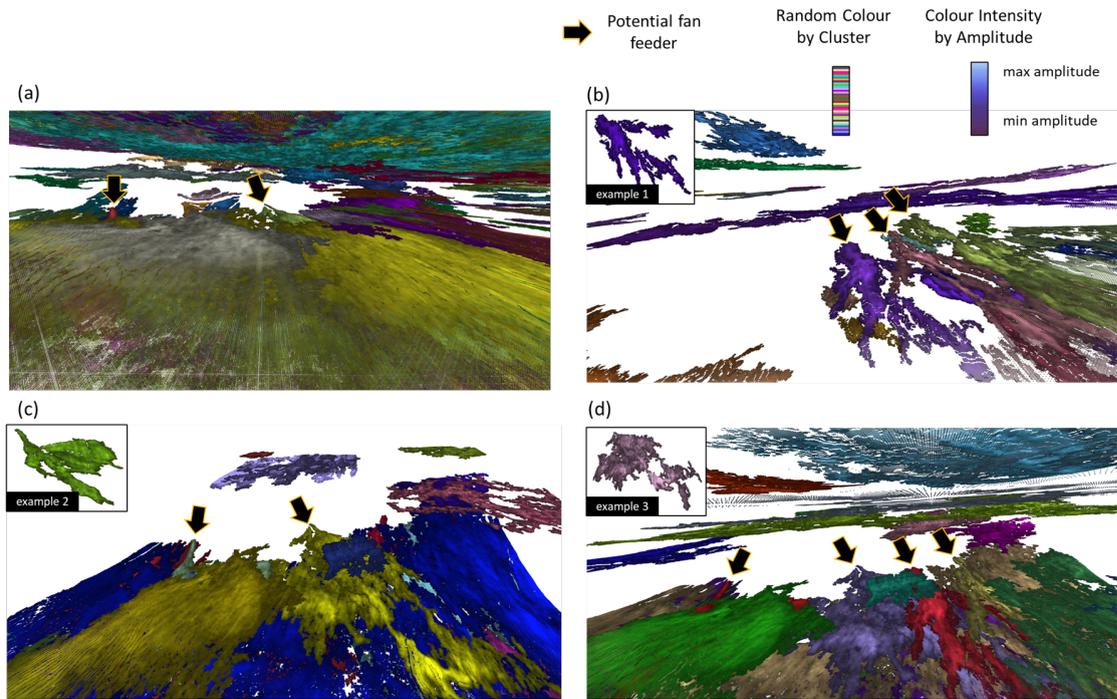


FIGURE 4.19: Result comparison of the retrieved objects inputting three different fans as an example.

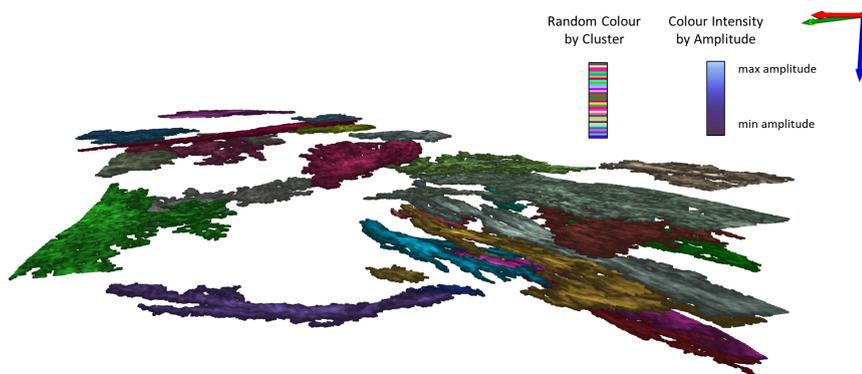


FIGURE 4.20: Results of the 30-closest objects retrieved from the elongated fan. Zoom out from Figure 4.19(d)

In addition to the closest objects, the method could return the similarity distance between objects, which can be seen as a proxy for the uncertainty of the detection. The precision of the similarity distance between objects depends on the ability of the extracted features to fully characterise objects and should therefore be used with caution. However, the similarity distance between objects was not directly analysed in this work.

4.5 Conclusion and discussion

The results presented in this chapter have shown interesting results and potential for the characterisation and detection of fans in seismic data. First, a feature extraction approach allows the extraction of specific characteristics of each object in the segmented seismic. Then, a closest object retrieval approach allows retrieving n objects most similar to an example input by the interpreter. The method allows the rapid and successful identification of certain geobodies; we applied it to the *validation dataset* and could find the three manually interpreted fans and potential other fans that were not interpreted.

The robustness of these results needs to be further validated and tested on data that were not considered during the development of these methods. Testing on completely unseen data allows to produce an unbiased estimate of the final model performance. In Chapter 5, we test the adaptability to unseen data and other types of geobodies and are presented in Chapter 5.

Another interesting perspective this workflow offers is the sensitivity analysis of uncertain parameters for geobody detection. The principle of sensitivity analysis is to vary the parameters of the method and to compare the different scenarios generated. These scenarios can be analysed according to the accuracy of detection for each geobody but also in terms of variations in the delineation of the extracted objects. The main parameters that could potentially lead to a variety of interpretations are: (i) the seismic attributes used to extract the point cloud, different seismic attributes underline different aspects of the seismic data, and this could lead to different definitions of the seismic objects; (ii) the segmentation parameters (*cutoff*, *eps*, *MinPts*), changing the segmentation parameters will impact the reach and connectivity of the clusters and therefore the segmentation level; and (iii) the feature selected to characterise the seismic objects, leading to different similarity distances between objects.

The feature extraction is a sensitive part of the methodology and defines the success of the detection, but the results presented in this chapter can be improved. The features presented in this chapter are those that we consider the most relevant for characterising seismic objects while remaining general, but other features deserve to be explored.

Another approach to feature extraction that may have merit is to use an auto-encoder network to extract features from a point cloud object. An auto-encoder compress (encode) data into fewer dimensions and then decompress (decode) it to reconstruct the object with minimal errors. Thus, auto-encoder has been used for dimensionality reduction [Alsenan et al., 2020, Wang et al., 2016] and feature extractions, using the reduced dimensions (latent-space). The principle of the auto-encoder approach would be to train a network to receive a seismic object (point cloud) as input and reconstruct the object as the output of the encoder-decoder sequence (figure 4.21). Then, the encoder part would be used to extract a dimensionality-reduced representation of the object and use that representation as features characterising the object. The complexity will be to build a network that directly handles point clouds as inputs and outputs, with point clouds of different sizes, probably using a graph-based neural network [Elbaz et al., 2017, Yang et al., 2018]. One drawback of this approach is the loss of control over the features and their meaning, but a properly trained auto-encoder should significantly improve the accuracy of object characterisation.

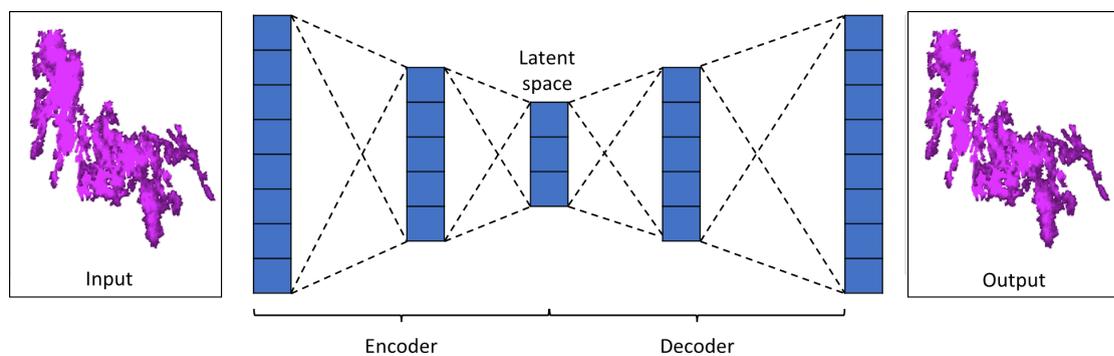


FIGURE 4.21: Encoder decoder network structure. The encoder part of a trained network could be used for feature extraction, using the latent space representation as features.

Chapter 5

Geobody detection in 3D seismic - Application to real case studies

5.1 Introduction

In chapters 3 and 4, we introduced a novel workflow to facilitate the detection of geobodies in 3D seismic data. In this chapter, we will test the methods on two real seismic datasets to assess the accuracy and robustness of geobody detection on real data. In chapters 3 and 4, we validated each method of the workflow on a *validation dataset* that we also used for developing the method. Testing on completely unseen data allows to produce an unbiased estimate of the final model performance, while, for the *validation dataset*, the fact of knowing the data when designing the methodology introduced a bias in the evaluation. Therefore, we now test the method on the unseen data at the time of the development of the methods. The two tests consist of the following:

- (i) Apply our workflow to the complete seismic survey from the North Falkland Basin (NFB), aiming to discover new fans over and above the three manually interpreted and presented in chapter 3 -North Sea Lion, Sea Lion and Casper fan-; and
- (ii) Apply our workflow to another seismic survey from the Falkland Plateau Basin (FPB), aiming to retrieve a different type of geobody: sill intrusions, thus assessing

the robustness and the adaptability of the method to different geologic settings, seismic quality and types of geobodies.

5.2 Application to turbidite fan detection in the complete North Falkland Basin dataset

5.2.1 Data and methodology

To carry out this research, we have at our disposal a significant, high-quality dataset from the North Falkland Basin provided by the Falkland Island Government. This dataset contains a modern 3D seismic reflection volume, core analysis, wireline logs, geochemical analysis and other diverse types of data. The 3D seismic covers 4500 km^2 in the NE part of the Northern Rift Basin and was acquired between 2004 and 2011 and processed together in 2011, resulting in interpretable seismic images of good quality down to $5s$ TWT with a $12.5m$ recording interval [Bunt, 2015]. The dataset displays several Cretaceous-aged fan bodies deposited into the Eastern Graben of the basin. Some turbidite fans have been manually interpreted from the 3D seismic data, and core data have been extracted from exploration and appraisal wells, with wells crossing the fans in several places [Bunt, 2015, Dodd et al., 2019, Richards et al., 2006].

The North Falkland Basin is a failed rift system [Richards et al., 1996]. The northern part of the basin consists of two half-grabens separated by a ridge. The fans were deposited in early post-rift via canyons from the eastern basin margin. The fans are either proven petroleum reservoirs or undrilled prospects with significant in-place volumes [MacAulay, 2015].

We investigate the ability of our point-cloud geobody detection workflow: (i) to operate on the typical size of seismic data required for an exploration campaign and (ii) to assist in the detection of fans. We presented results on the *validation dataset* from that seismic data in Chapter 3 and 4, which is a cropped cube from a larger 3d survey (figure 5.1(a)). The complete survey has about eight times the area of the *validation dataset*.

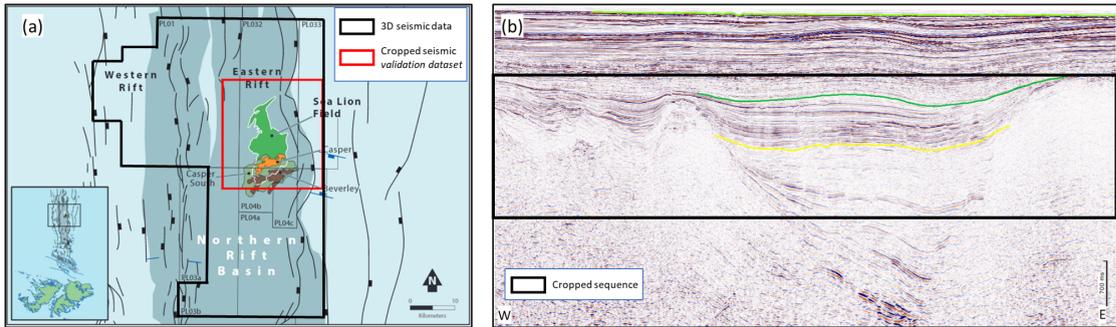


FIGURE 5.1: Presentation of the NFB seismic dataset: (a) map view and localisation of the complete seismic data and the cropped *validation dataset*, modified from Bunt [2015]; and (b) W-E section through the seismic dataset, cutting through the Eastern and Western Rift; the cropped sequence contains the fans and the results presented here are produced from that cropped sequence of the seismic data.

We check how the method behaves on the entire NFB seismic dataset. We will first evaluate the results by visually evaluating the quality of the segmentation and then by assessing the ability of our approach to detect new fans. We cut the seismic at the level of the stratigraphic sequences that contain fans, setting aside the shallow (above 1600ms TWT) and deep parts (above 3600ms TWT) that are not of interest for fan exploration (figure 5.1(b)). The *test dataset* for the NFB consists of a seismic cube of 3150 inlines, 5032 crosslines and 501 timeslices.

Processing step	time
Extrema extraction (maxima)	13min 12s
Normalize and filter amplitude	9s
Compute and filter coherence	8min 41s
Segmentation (DBSCAN)	9min 8s
Total	31 min 10s

TABLE 5.1: Processing times for the segmentation of the NFB dataset. Laptop with an Intel Core i9-10885h CPU processor and 32GB of RAM.

5.2.2 Results

5.2.2.1 Seismic segmentation

We perform the segmentation of the seismic using the methods introduced in Chapter 3: extraction of a point cloud and then segmentation of this point cloud with DBSCAN. For the point cloud extraction, we choose cut-off values of amplitude 0.4 and coherence 0.95. The point cloud extraction takes about 22 minutes to compute – including the computing of the coherence attribute – (Table 5.1). The total studied NFB seismic data consists of 7,941,250,800 voxels, which the point cloud extraction reduces to 113,738,264 points, giving us a reduction factor of 98.5% from voxels to points. Next, we perform the segmentation of the point cloud with DBSCAN using an *eps* of 2, a *MinPts* of 8 and no *Zexg*. The segmentation takes about 9 minutes to run (Table 5.1). The segmented seismic is composed of 8,725,973 noise points and 574,632 segments. Most of these segments are very small; filtering out the segments with less than 10,000 points reduces the number of segments to 544 (figure 5.2). The manually interpreted fans are composed of more than 100,000 points on average. Thus, filtering on objects with less than 10,000 points should not suppress fans unless the fan is over-segmented into small parts. The 544 clusters constitute the database of seismic objects, and the objective is to extract the fans among all these objects.

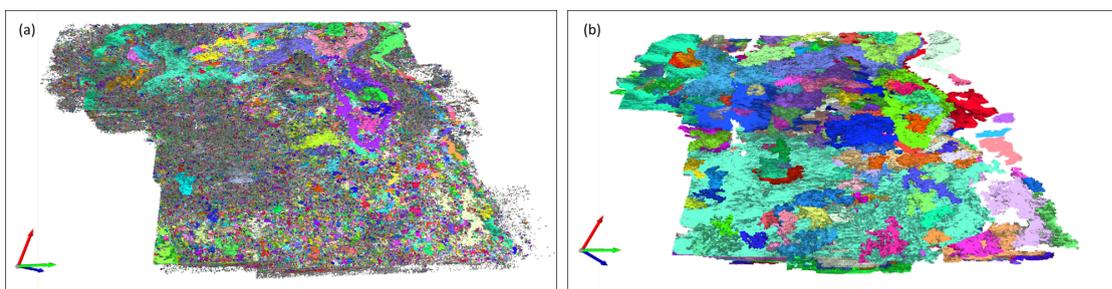


FIGURE 5.2: The segmented seismic of the NFB dataset is composed of 574,632 segments (a); filtering on small clusters reduces it to 544 segments (b).

The first observations from the segmentation results are:

- (i) The full cube seismic segmentation is fast, even for a relatively big seismic cube. The total processing time is about half an hour for the seismic point cloud extraction

and the segmentation, which is extremely fast for a seismic of this size (3150 x 5032 x 501) compared to other full seismic segmentation methods.

- (ii) The segmentation adapts easily to the entire seismic survey. By applying similar parameters as for the *validation dataset*, the results obtained are of good quality for the whole survey.
- (iii) Overall, the segmentation accurately delineates many objects in the seismic data and is of good quality with limited over-segmentation and under-segmentation. Thus, point cloud segmentation makes it possible to have a nicely rendered 3D visualization of the seismic survey and to start analysing the distribution of the objects which compose it. For example, in Figure 5.3, we observe the two separated half-grabens and the infill objects of various sizes and extends.

Some errors and variations in the delineation of geological objects can be observed when exploring the segmentation results. For example, in Figure 5.4, a fault splits the end of an object in two, resulting in a delineation into two different segments, while the spatial continuity of the amplitude values clearly indicates that it is the same event. In such situations, the merge optimisation method proposed in Figure 3.31 could improve the segmentation and ideally correct such errors. It is also important to note the differences in delineation caused by the choices of segmentation parameters. Comparing the Sea Lion delineation with the result of the *validation dataset* and the NFB dataset, differences are noticeable, although the differences in the segmentation parameters are small (figure 5.5). Firstly, we observe an additional body at the top right corner of the delineation produced in Figure 5.5(b.iii). This additional body is probably a small fan so close to the Sea Lion fan that they are segmented together. Secondly, the upper left part of the Sea Lion is more complete in Figure 5.5(b.i) than in Figure 5.5(a.i), probably due to lower amplitude and coherence cut-offs. Finally, the low, channelised termination of the Sea Lion is missing in Figure 5.5(b.ii), probably due to a higher point-connectivity threshold in the DBSCAN parameters. Thus the delineation of 5.5(b) is both more complete in some aspects (circle (i)) and less complete in others(circle (ii)). However, all these observations could be interesting from a reservoir connectivity perspective; for instance,

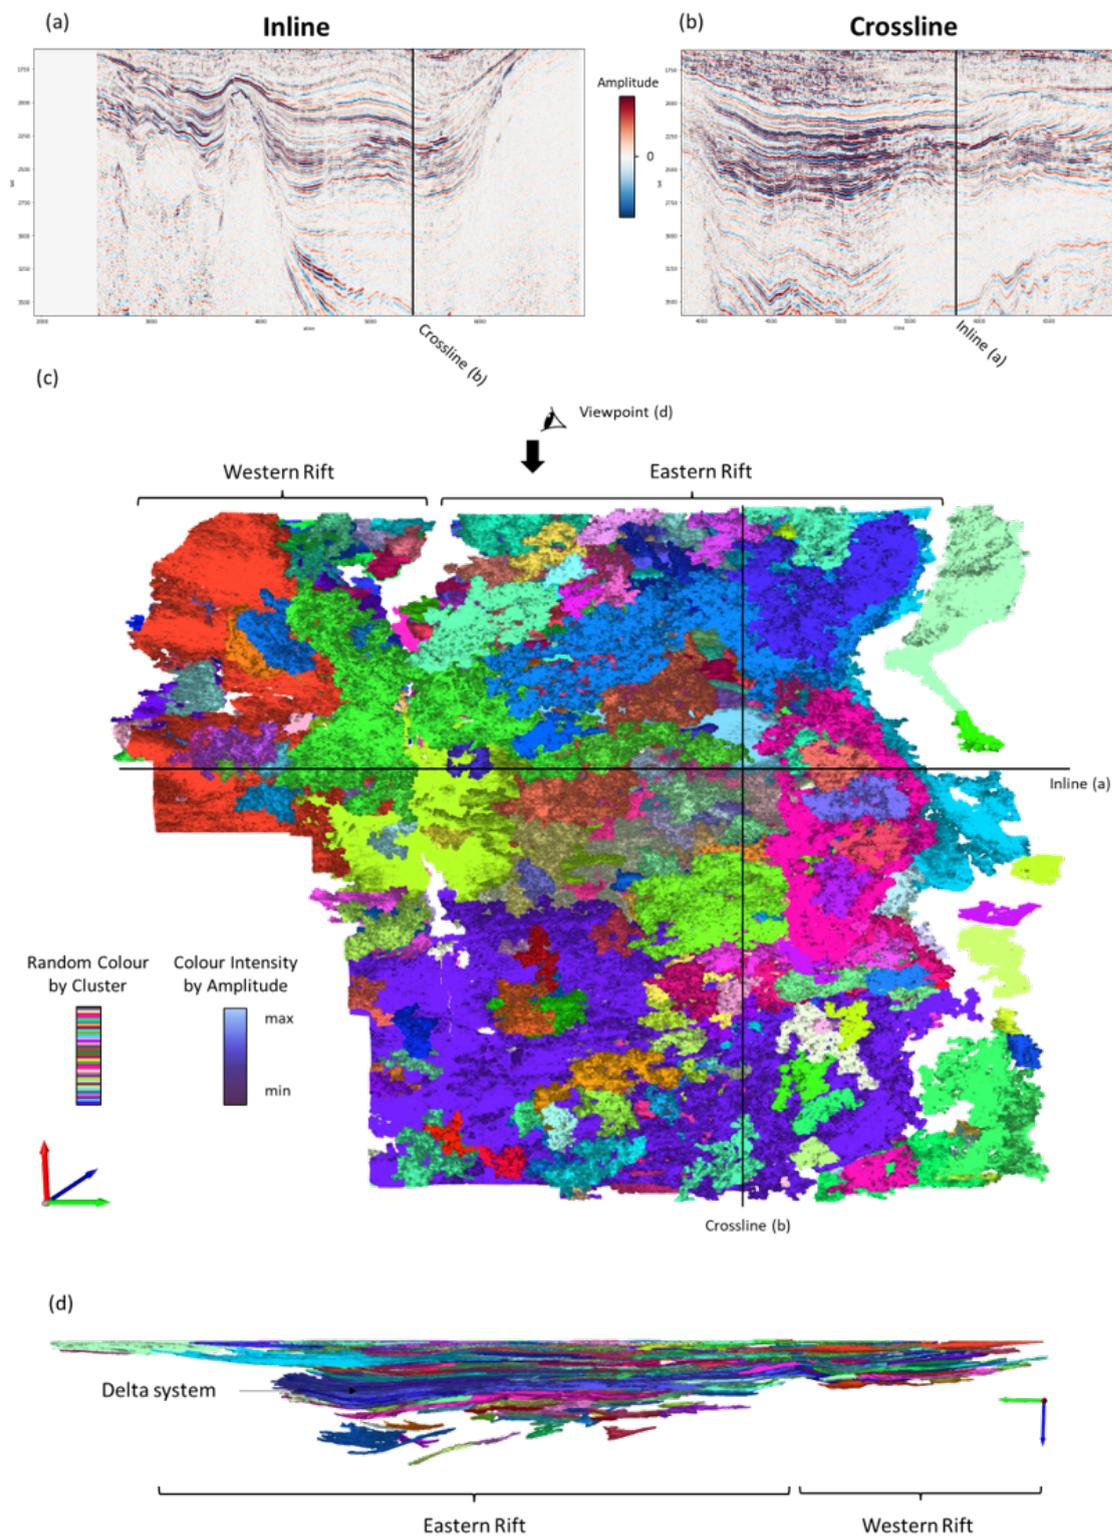


FIGURE 5.3: Results of the seismic segmentation of the entire NFB dataset. (a) and (b) are two seismic sections of the seismic data input to the segmentation, and (c) and (d) are two orientation views of the output of the segmentation.

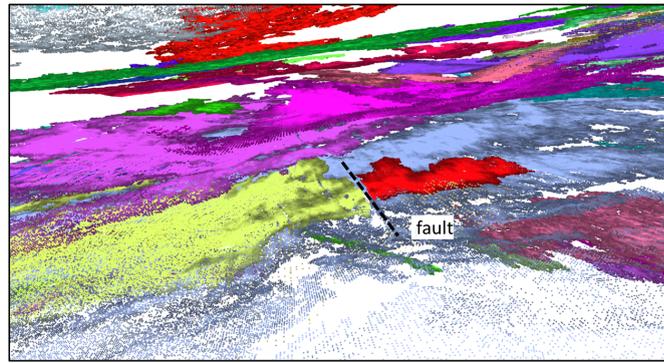


FIGURE 5.4: Fault splitting an object and resulting in a segmentation error.

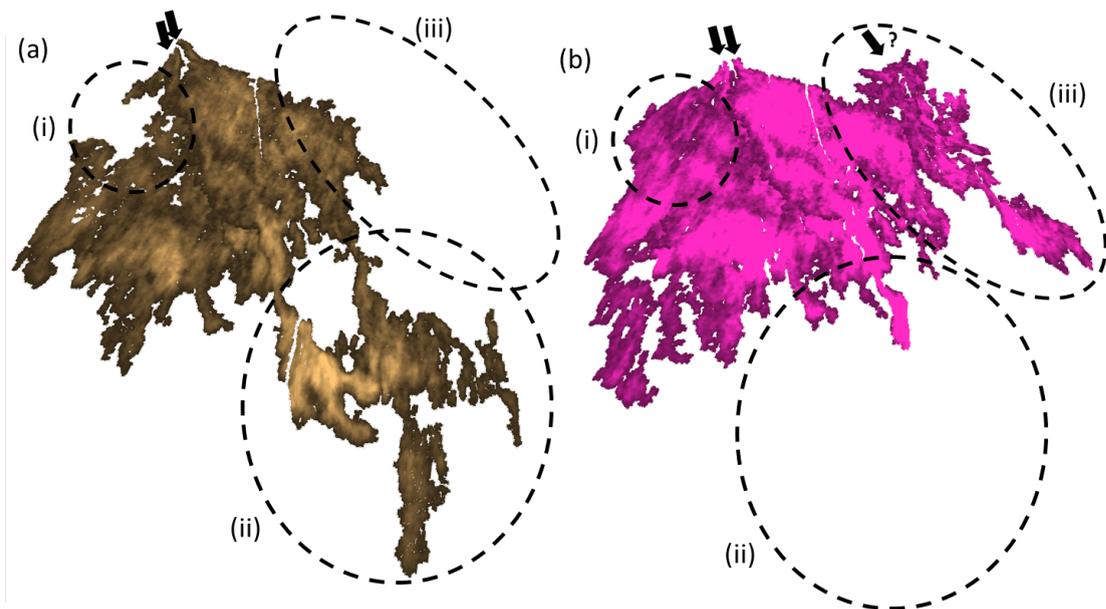


FIGURE 5.5: Comparison between the Sea Lion fan delineated from the segmentation of the *validation dataset* (a) and of the entire NFB seismic survey (b). The dashed circles highlight the main differences between the two segmentation results.

we can ask ourselves whether the channelised termination (circle (ii)) is connected to the rest of the fans in terms of a reservoir or if the reservoir is compartmentalised.

The analysis of the largest cluster in the segmentation is interesting and can be seen in the context of the basin analysis presented by Bunt [2015]. Figure 5.6 shows the depositional patterns visualised with the point cloud segmentation and the spectral decomposition techniques. The first observation is the major delta system prograding from the north. The point cloud representation allows the quick 3-dimensional screening of the different depositional layers of the delta and their connectivity. Other depositional events can be observed and are similar to what Bunt describes: viewed from below, a major fan

can be seen on the lower SW corner (event **B**); and, from the top view, another fan can be seen (event **A**), showing the change in orientation of the primary sediment inputs, now coming from the SE. An ideal delineation would have separated these different elements into several segments. A segmentation optimisation process (see Chapter 3 – Optimisation of the Segmentation) might improve this result. Still, it will be common for large depositional events, such as this delta, to be so closely connected to surrounding events deposited in the same time interval that they will be segmented together.

5.2.2.2 Fan detection

The first observation we made for fan detection was to check the delineation of the three known fans from the *validation dataset*: the Sea Lion fan, the North Sea Lion Fan and the Casper fan. The three fans are easily recognised by looking through the segmentation results and are stratigraphically above the deltaic deposits (figure 5.7).

The object detection methodology is applied to fan detection in this first test and uses the methods introduced in Chapter 4 to retrieve the 100 objects most similar to a fan example. The features selected for fan characterisation are the same as those used for the validation dataset: the mean amplitude, the mean coherence, the mean Zeboudj's distance, the number of points, the contour ratio, the linearity and the slope. The feature extraction takes about 4 minutes to characterise all the objects. The object used as an example for the n-closest object retrieval (with Locality Sensitivity Hashing (see Chapter 4) is the Sea Lion fan, and we are looking for fans among the 594 objects that constitute the database. We ask for the 100 closest objects, and the retrieval of objects takes a few seconds to run.

The search for similar objects in a segmented point cloud seismic data facilitates the detection of fans. Figure 5.8(a) displays the results of the 100 closest objects and a selection of fourteen objects interpreted as potential fans. The fourteen objects were identified manually by screening through the 100 object results. The 2D image limits the understanding of these point clouds, but these observations are made on the basis of interactive 3D views that allow the characteristics of these fans to be better distinguished. First of all, we find the three fans manually interpreted -objects number 9, 10 and 11 in

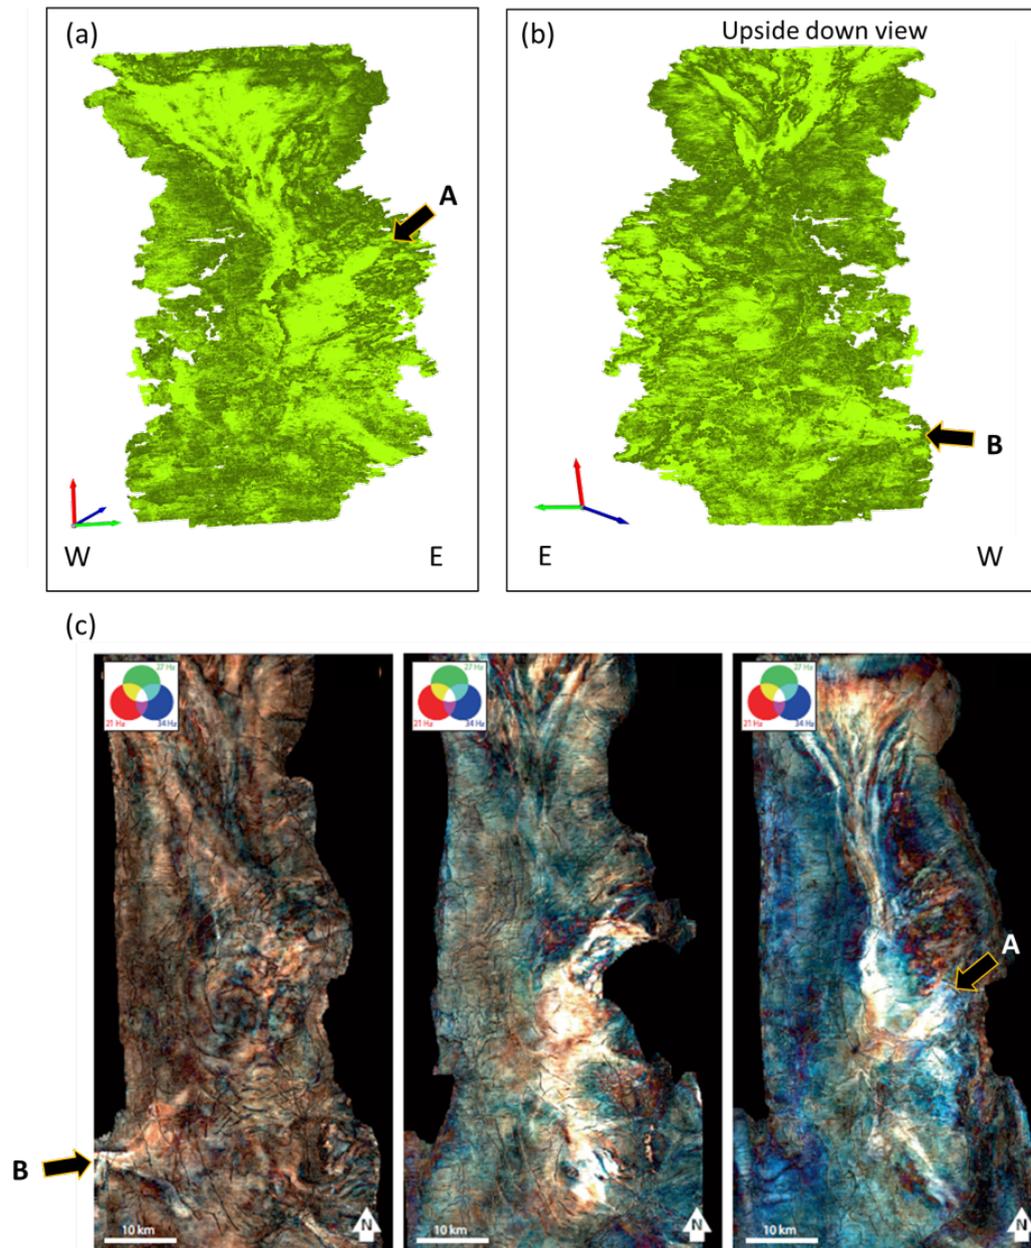


FIGURE 5.6: Delta system observed both from point cloud segmentation and spectral decomposition visualisations (a) view from the top of the largest segment, (b) view upside down of the largest segment, and (c) Regional spectral decomposition horizon slices through the Northern Rift Basin within the NFB from deeper (left) to shallower (right). Figure (c) from Bunt [2015].

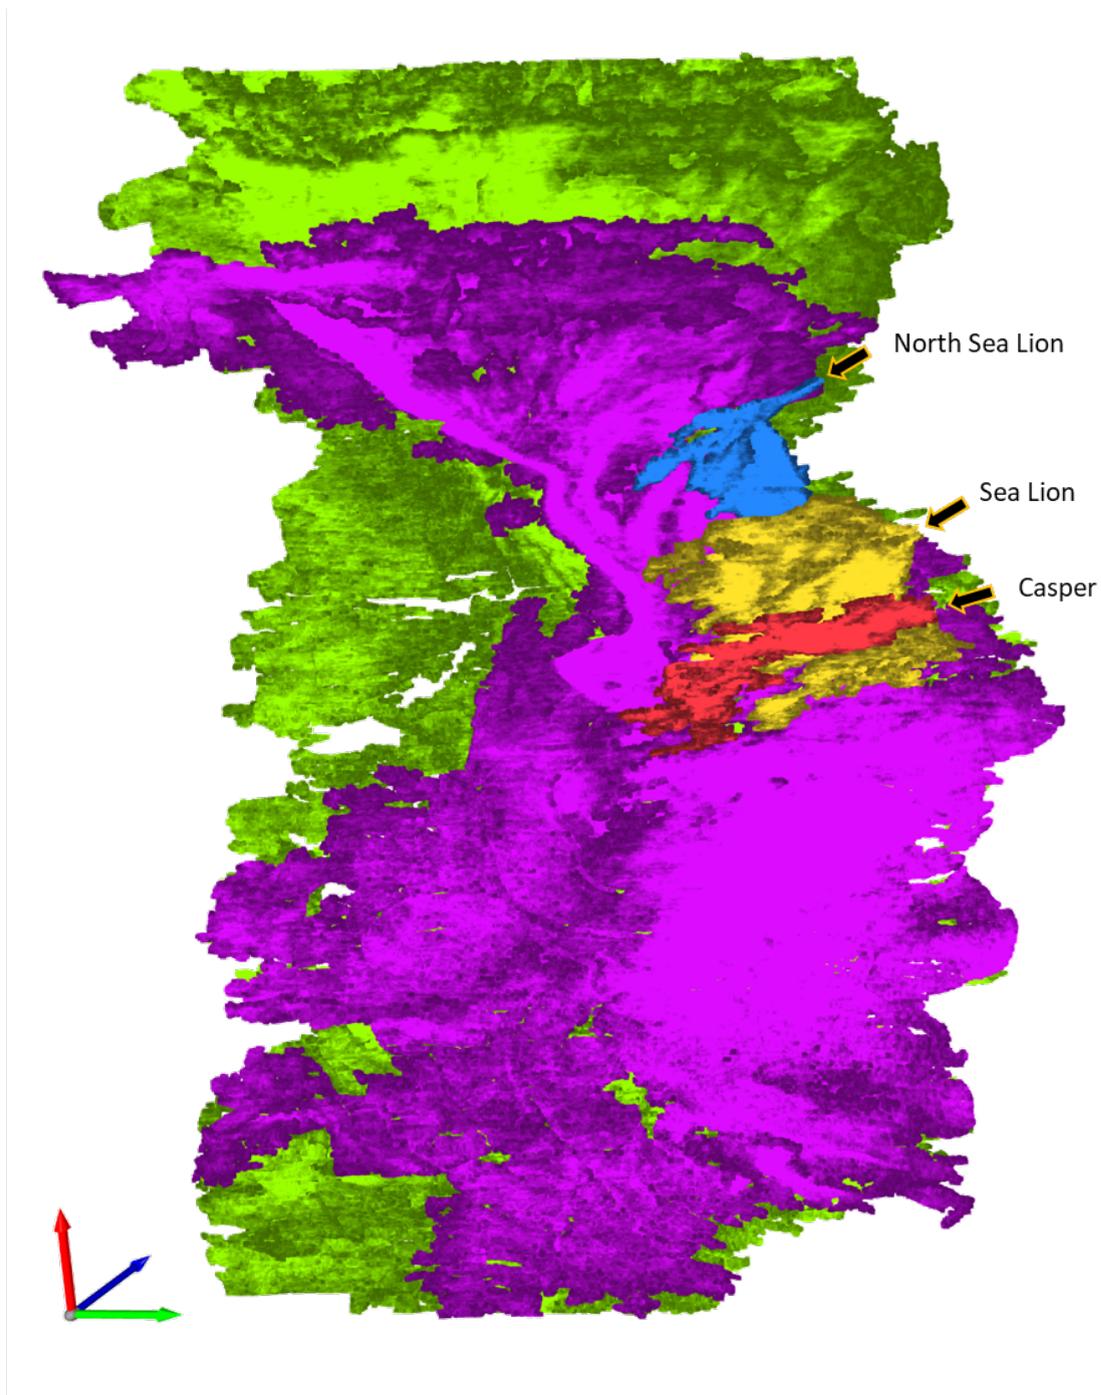


FIGURE 5.7: Delineation of the three known fans displayed above the deltaic deposits (the two largest clusters of the segmentation).

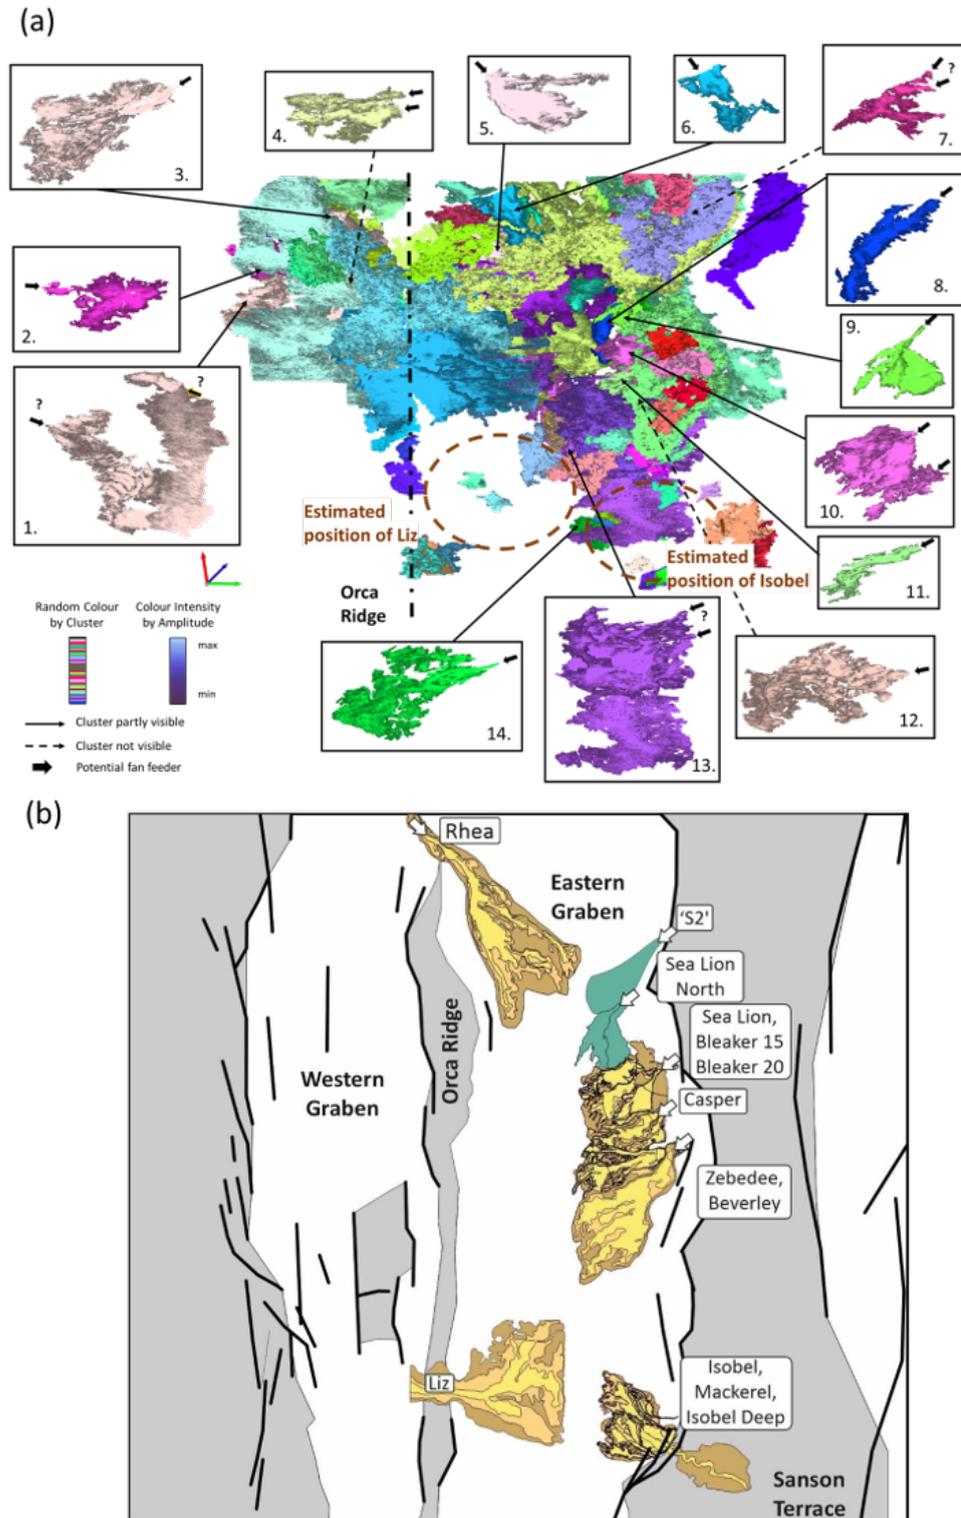


FIGURE 5.8: (a) Results of the 100 closest objects to the Sea Lion fan by LSH and analysis of the potential fans observed; (b) Position of the main fans interpreted in the NFB -used as ground-truth detection-. Figure from Dodd [Unpublished].

Figure 5.8(a)-. Ten potential fans are located in the Eastern Graben, with sediment inputs mainly coming from the eastern margin (objects 5 to 14), and four potential fans are

located in the North of the Western Graben (objects 1 to 4), with sediment inputs either from the east (intragaben high - Orca Ridge) or the western margin. By comparing these potential fans with the fan map drawn up by Dodd [Unpublished] (figure 5.8(b)), we can identify with varying degrees of confidence five of the seven interpreted fans: Rhea (object 5 in figure 5.8(a)), Sea Lion North (object 9), Sea Lion (object 10), Casper (object 11) and Zebedee and Beverley (northern part of object 13). But two interpreted fans are missing: Liz and Isobel. Liz is not present among the retrieved objects (figure 5.8(a)). For Isobel, some objects are present at the same location, but we cannot identify the channelled patterns that would give confidence in the identification (figure 5.8(a)). The faults running through Isobel's feeder could be the reason for this absence, leading to the coherence cut-off filtering the points at that location. All of these interpretations would benefit from further validations than based on a single map, but this already provides a first positive validation. Again, the segmentation could be improved; for example, object 13 delimits a fan but also includes a horizon on its southern side, which complicates the detection by fan-specific features.

5.2.3 Assessment of the application to the NFB dataset

This first application successfully tests using our segmented point cloud and similarity-based object search methodologies on a basin-wide seismic survey to find fans. We had no prior knowledge of some of these fans at the time when the method was developed, thus removing some of the bias from the evaluations. The major benefits are the speed of the method and the interactivity offered when exploring seismic data for geobodies. The method allows the seismic to be interrogated in a different way while remaining in control: the seismic interpreter navigates through the point cloud to gain a more global understanding of different geological structures and selects potential targets for further analysis more quickly and easily. The aim is not to replace the interpreter but to provide them with an additional tool to make better-informed decisions.

5.3 Application to sill detection in the Falkland Plateau dataset

The second application on the FPB allows for testing on a new data set and a new object type: a different seismic dataset and searching for sill intrusions. The application on the NFB allowed us to test the method on unknown parts of a seismic survey. However, the seismic data was still the same as the *validation dataset*. A different seismic dataset allows us to remove any bias and evaluate the robustness of various seismic qualities in terms of processing and resolution of the seismic data. In addition, we use a different type of geobody to assess the flexibility of the method to other shapes and geometric descriptors.

5.3.1 Data and methodology

For this second application, we have at our disposal a 3D seismic survey (FISA survey) from the FPB, located east of the Falkland Islands (figure 5.9) which shows evidence of numerous Early Cretaceous dyke and sill intrusions [Richards et al., 2013, Stanca et al., 2022]. The 3D seismic cube covers 5500 km^2 with high-resolution, interpretable data 9s TWT, with a 12.5 m recording interval and a sampling rate of 2ms [Stanca et al., 2022] (figure 5.11).

An important series of intrusive magmatic dykes and sills intruded the sedimentary succession filling the FPB. The igneous intrusions were dated from two phase of magmatic activity, one from the Jurassic and another from the Cretaceous [Stone et al., 2008]. Stanca et al. [2022] have mapped the sills across the basin and characterises them as mainly saucer-shaped bodies with high amplitudes and stepped geometries (figure 5.10).

We investigate whether our method can extract the dyke and sill network from a single sill example. We cropped the seismic sequence of interest from the FISA seismic cube where the magmatic bodies intrude the sedimentary layers (figure 5.11). The cropped seismic volume includes 2636 inlines, 2251 crosslines and 1000 timeslices. We will first

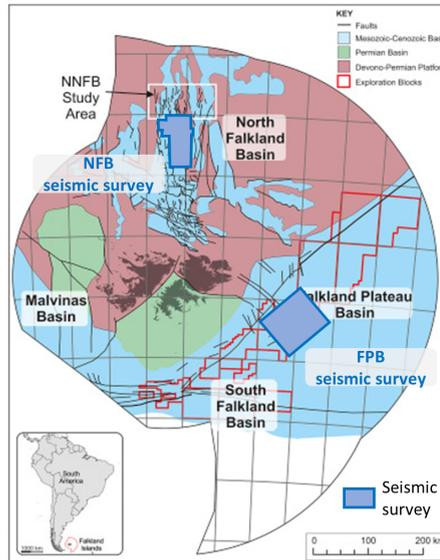


FIGURE 5.9: Location of the FISA seismic survey in the NFB.

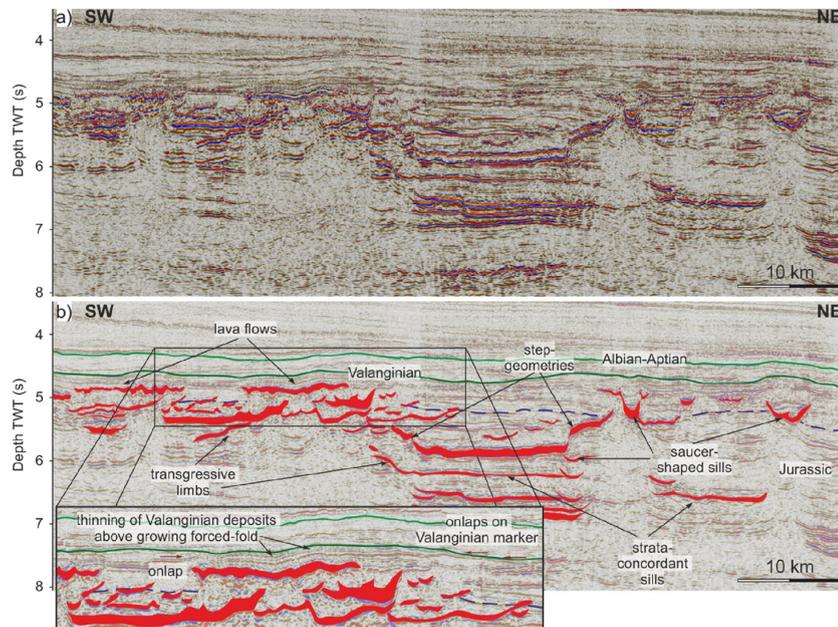


FIGURE 5.10: Sills and lava flow distribution and associated forced-folds, mapped by Stanca et al.. Figure from Stanca et al. [2022].

evaluate the results by looking at the quality of the segmentation and then in relation to sill detection.

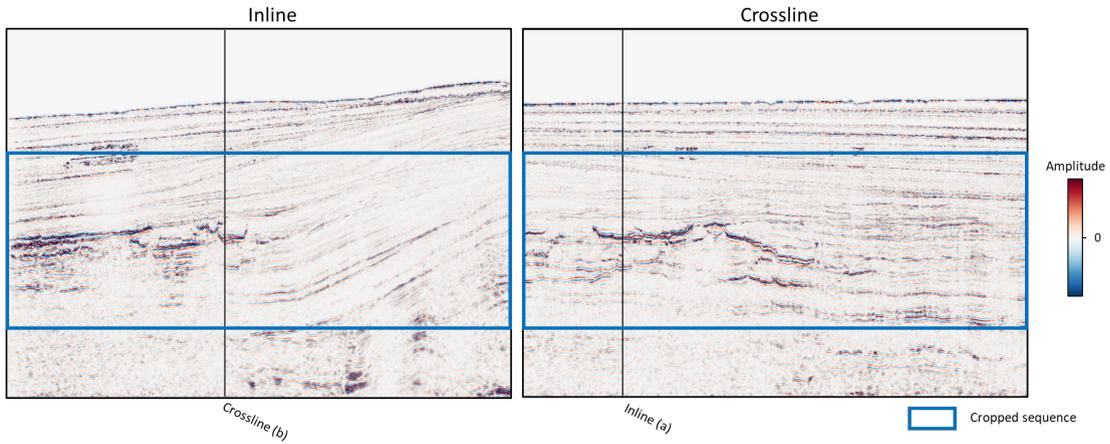


FIGURE 5.11: Cropped sequence from the FISA seismic cube - FPB -, where the magmatic bodies intrude the sedimentary layers.

Processing step	time
Extrema extraction (maxima)	36min 39s
Normalize and filter amplitude	9s
Compute and filter coherence	34min 30s
Segmentation (DBSCAN)	18min 15s
Total	1h 29min 33s

TABLE 5.2: Processing times for the segmentation of the FISA dataset, FPB. Laptop with an Intel Core i9-10885h CPU processor and 32GB of RAM.

5.3.2 Results

5.3.2.1 Seismic segmentation

For the segmentation of this new seismic data, we follow the same methodology: extraction of a point cloud and segmentation with DBSCAN. The point cloud extraction uses cut-off values close to those used for the first seismic survey, with a cut-off amplitude of 0.35 and coherence of 0.95. The point cloud extraction takes about 1 hour and 10 minutes to compute (table 5.2). The total studied seismic consists of 5,933,636,000 voxels, and the point cloud extraction reduces it to 208,007,963 points, which gives us a reduction factor of 96.5% from voxels to points. Similarly, we perform the segmentation

with the same DBSCAN parameters using an eps of 2, a $MinPts$ of 8 and no $Zexg$. The segmentation takes about 18 minutes to run (table 5.2). The segmented seismic is composed of 36,331,158 noise points and 1,594,216 segments. But, again, most of these segments are small and filtering out segments with less than 10,000 points reduces the number of segments to 1263. Igneous intrusions can vary greatly in size, but for the purposes of this application, it is reasonable to keep only intrusions consisting of more than 10,000 points.

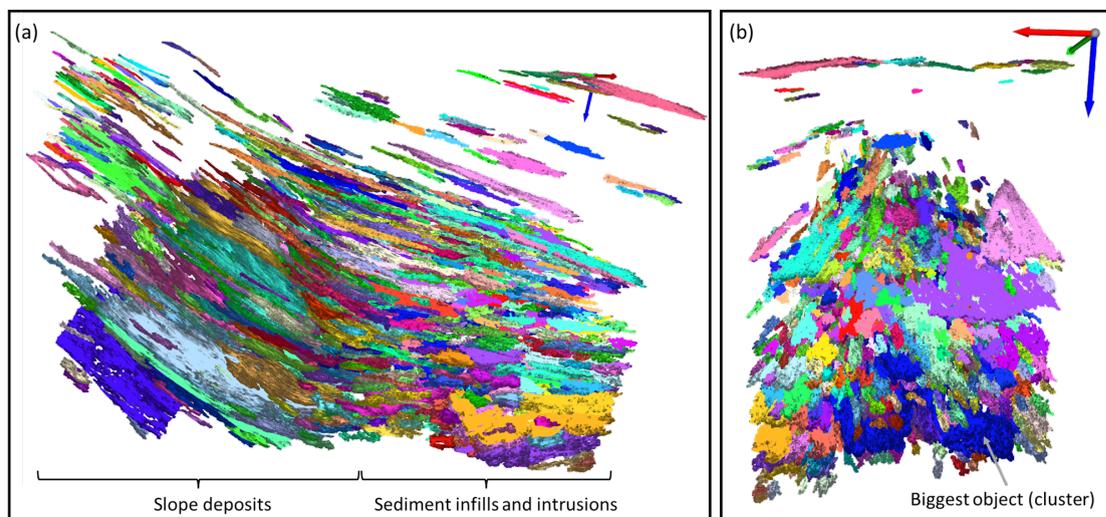


FIGURE 5.12: Results of the point cloud segmentation of the FPB dataset. (a) and (b) display two different view angles of the 1263 objects constituting the object database. A 4x vertical exaggeration is applied to the view.

Figure 5.12 displays the results of the segmentation with 1263 segments over 10,000 points. The upper part is quite sparse as most of the elements have been filtered out due to more of a chaotic seismic reflection. In Figure 5.12.a, on the right, the segmentation delineates deposits along the slope of the basin and on the left, the sedimentary infills and magmatic intrusions. This is consistent with Duarte's observation that the intrusions are located in the south-eastern part of the 3D cube. Moreover, the larger cluster delineates the intrusions that correspond to the central sill interpreted (figure 5.13). This sill is thought to act as a central feeder [Duarte, 2019].

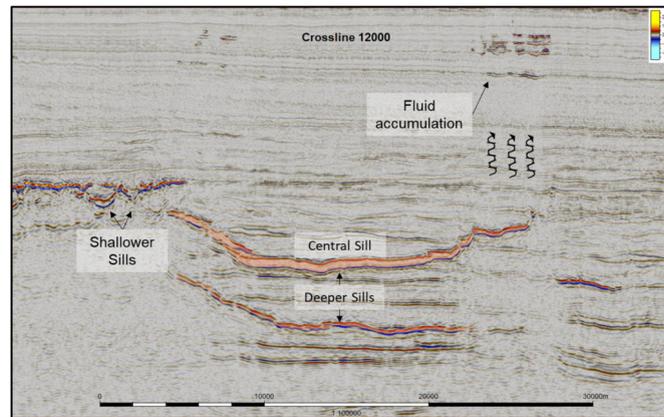


FIGURE 5.13: Vertical seismic section crossing the central sill reflection, FPB. Figure from Duarte [2019].

5.3.2.2 Sill detection

The object detection methodology is oriented towards sill detection, aiming to retrieve many sills from a single sill example. So, the first thing to do is to interpret an example of a sill from which our approach will search for similar objects. The example we picked is a sill from the shallow part of the seismic data that we could easily characterise by its saucer-shaped body and high amplitude reflectivity (figure 5.14). This sill is relatively small in size (55637 points) compared to the larger, deeper sills.

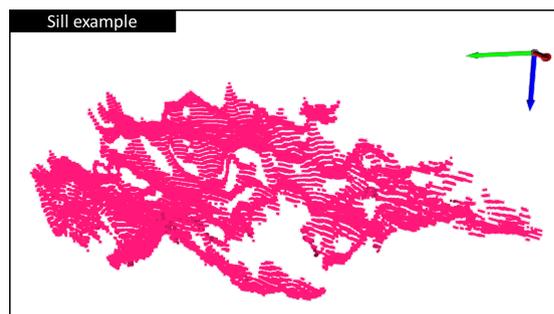


FIGURE 5.14: 3D visualisation of the shallow sill provided as an example.

Similar to our fan workflow, we characterise the set of 1263 objects of the segmented seismic data with features that best characterise sills. The features selected are the mean amplitude, the mean coherence, the mean Zeboudj's distance, the number of points, the contour ratio, the linearity, the sphericity, the planarity and the point elevation distribution. Compared with the fan recognition, we add planarity and point elevation distribution as features because we consider these to be essential for characterising the

vertical elevation of the tips of sills. The feature extraction takes about 9 minutes to run for all 1263 objects. Then, we ask for the 100 closest objects to the sill example using LSH, and retrieving objects take a few seconds to run.

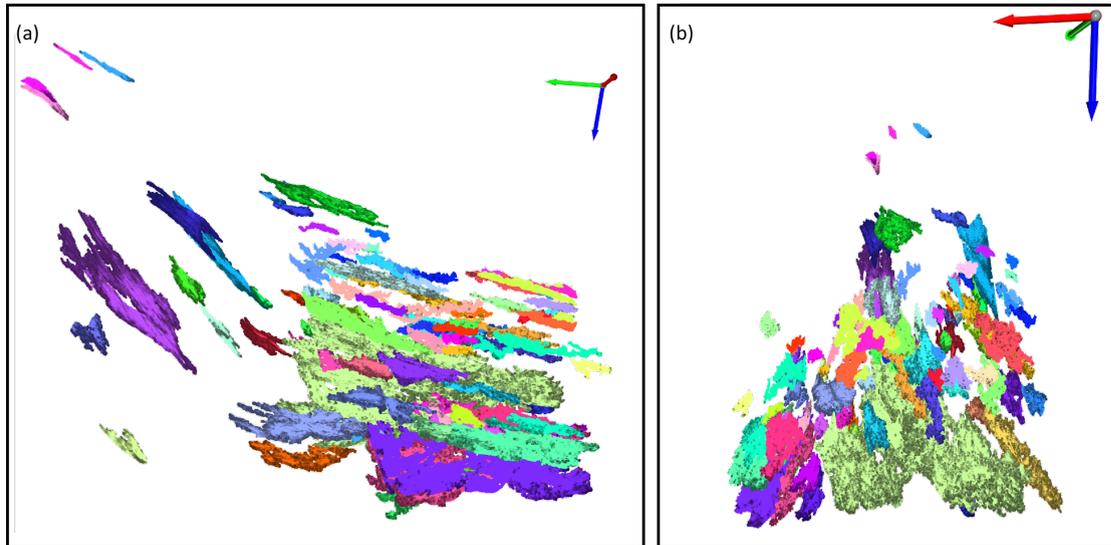


FIGURE 5.15: Sill detection results - 100 closest objects to a sill example. (a) and (b) display two different view angles of the same 100 objects. A 4x vertical exaggeration is applied to the view.

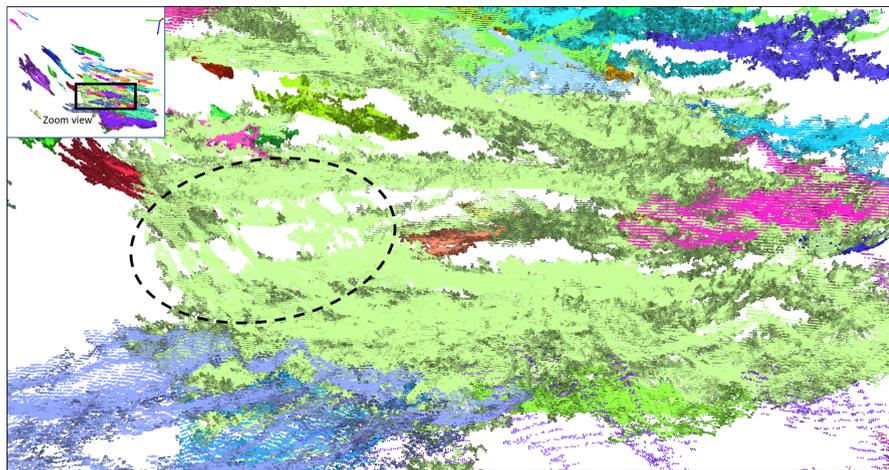


FIGURE 5.16: Focus on the main object retrieved, segmented point cloud showing the central sill and the connections of a network of intrusions with dykes that vertically link sills (green segment).

The extraction results are presented in Figure 5.15 and are aligned to the map of intrusion interpreted by Duarte. First of all, if we look at the overall distribution of the bodies obtained, the vast majority is indeed within the south-eastern extremity of the volume, as in the manual interpretations. Then, if we look in more detail at the main segment (figure 5.16), we can see that it is composed of sills (horizontal sheets) connected by dykes

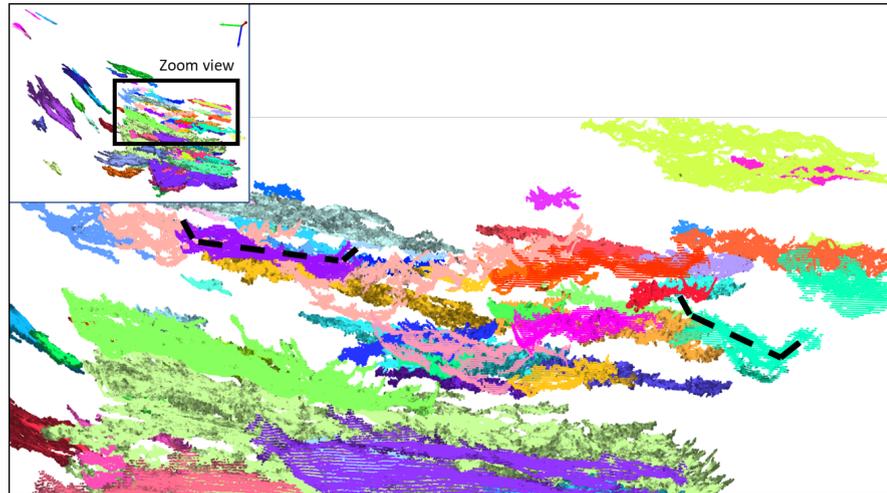


FIGURE 5.17: Focus on the objects retrieved at shallow depths, segmented point cloud showing small sills with saucer-shaped bodies. The dashed lines highlight the saucer shape of small sills.

(inclined sheets - dashed circle in figure 5.16). This main segment would correspond to the central sill attached to the deeper sill, observed in Figure 5.13, and would be joined by dykes. Finally, looking at the upper sections, we observe the shallower sills with their characteristic saucer shape (figure 5.17).

5.3.3 Assessment of the application to the FPB dataset

The results of this second application confirm (i) the efficiency of the point cloud segmentation to delineate a set of objects in a seismic volume, (ii) the possibility of finding some of the sills by similarity to an example and (iii) highlight the flexibility of the methodology for different seismic data and objects.

A significant advantage of point cloud segmentation is the ability to segment an element over multiple z-horizons. For example, the case of the central sill segmentation showed the possibility of grouping a network of sills and dykes as long as they are point-connected. This contrasts with a limitation of many seismic auto-trackers that do not allow tracking events on multiple elevations; with many auto-trackers, horizon surfaces are only picked following the local planar reflections.

5.4 Conclusion and discussion

These two applications demonstrate the value of the novel methodology on the data unknown at the time of development, thus producing an unbiased estimate of the performance of the methodology.

The method allows the rapid and successful identification of certain geobodies -sills and fans in these applications-. However, one could apply the same method to other geobodies, provided they are continuous types of bodies to allow segmentation by point connectivity. This adaptability is one of the advantages of an unsupervised approach, from the segmentation stage to the object identification stage.

The closest object retrieval approach significantly reduces the number of potential targets, making it more manageable for the interpreter to validate each object's identification precisely. Moreover, the interactivity and sparsity brought by the 3D visualisation of the point cloud make that identification significantly easier.

Finally, a level of control is left to the user. Depending on the choice of example type provided to the algorithm, the interpreter directs the search for the object type to its criteria. Again, the aim is not to replace the interpreter but to provide them with an additional tool to make faster and better-informed decisions.

In order to further evaluate the robustness of the segmentation results, it would be necessary to perform an application on data containing highly faulted areas or with a low signal-to-noise ratio.

To further evaluate the geobody search, an interesting test would be to take an object from another seismic as an example and search for similar objects in another seismic data or even an object synthetically modelled.

Chapter 6

Conclusion, discussion and future work

6.1 Conclusion

In this work, we introduced a new methodology to automate and facilitate the detection of geobodies in 3D seismic data. Geobody detection is essential to understand and accurately characterise the subsurface. However, faced with increasingly large seismic datasets, locating these subtle events is complex and time-consuming, leading to a need for automation.

The novel methodology allows, from a 3D seismic volume, to find items that look like a geobody of interest and to screen them (figure 6.1). The workflow consists of two main steps. The first step is to segment the set of objects present in the seismic volume; the segmentation is based on extracting a point cloud and clustering these points using the DBSCAN algorithm. The second step allows the extraction of n objects of interest according to an example provided by the interpreter. Thus, an interpreter can quickly query a seismic volume for certain geobodies without requiring any manual interpretation.

The uniqueness of this new approach is that the identification of geobodies is based on object recognition techniques that take advantage of the 3D geomorphological characteristics of the objects. Because all objects are delineated in the seismic point cloud in the first instance, the object recognition methods are able to work directly on the 3D shape

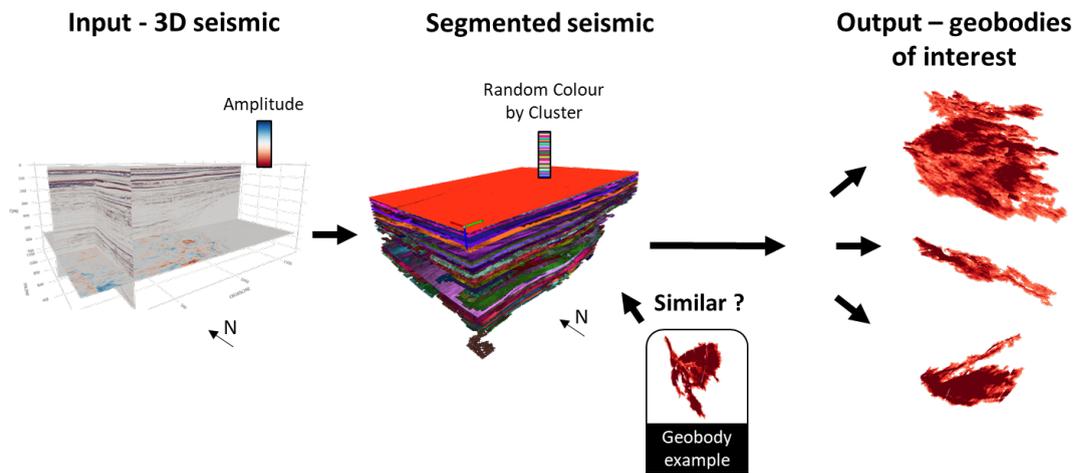


FIGURE 6.1: Summary of the methodology presented in this thesis, detection of geobodies from a 3D seismic data in a data-driven approach.

and amplitude reflection distributions of each object, which are fundamental criteria for the accurate characterisation of seismic objects.

The proposed methodology is based on unsupervised machine learning methods, so there is no prior training in a model required, which makes the methodology directly usable on any seismic data and adaptable to many types of geobodies.

Four main parameters control the seismic data segmentation: the amplitude cut-off, the coherence cut-off, the epsilon distance in DBSCAN and the minimum number of points in DBSCAN. Variations in the choice of these parameters allow the interpreter to vary the level of segmentation, from under-segmented to over-segmented. The choice is often a trade-off between the desire to create an accurate delineation of each structure and to preserve a maximum of objects.

In addition to the control of the segmentation, there is also the control of the object search, which the example provided by the interpreter orients. The choice of the shape and texture of objects to be retrieved from the seismic data allows the object search to be adjusted to the interpreter's desire.

From the different applications we have made, here are the key findings:

1. The visualization of the seismic in the form of a point cloud allows an understanding at a glance of the structures present in the seismic data. The sparse nature of

the point cloud allows seeing deep inside the seismic data without being restricted to sections of the volume. Moreover, the advances in graphics processing capabilities now allow the interactive display of point clouds with hundreds of millions of points. Therefore, a seismic point cloud approach may have benefits beyond the search for geobodies.

2. The segmentation of 3D seismic point clouds is extremely fast and scales well with the size of the seismic data. When applied to seismic data of several tens of GB in size, segmentation is achieved within an hour on a standard laptop computer.
3. Point cloud segmentation can be applied to most seismic data with little or no modification in the segmentation parameters and produce acceptable results.
4. The results applied to the detection of sills and fans show that it is possible, to some extent, to reduce the number of potential target objects to a subset of geobodies of interest. The characterisation of geobodies is a very complex task due to the diversity of geobody shapes, making it extremely difficult to hard-code the identification of geobodies using simple taxonomic rules. However, by extracting relevant geomorphologic features from each object and using methods such as Locality Sensitivity Hashing (LSH) to explore that feature space, it is possible to facilitate the identification of geobodies. The closest object retrieval approach allows the interpreter to extract similar objects from a unique example object, allowing the adjustment of the detection to the specific characteristics of an object.

6.2 Discussion

This discussion offers avenues for reflection on various aspects: (i) the scope of applications of the proposed methodology, (ii) the evaluation of the quality of a seismic interpretation, and (iii) opportunities for immersive visualisation.

In this thesis, different applications have been carried out in order to validate the proposed methodology but we can ask ourselves the question of the broader scope of application and limits of the point cloud seismic segmentation and geobody detection:

1. Point cloud segmentation provides a precise delineation of the majority of objects present in seismic, and its field of application can extend beyond the detection of geobodies, for instance, to produce a structural model with the extraction of the main seismic horizons. However, this full-volume segmentation cannot be used as a structural model. A structural model includes important interpretations embedded in the continuity of horizons and the delineation of faults that the point cloud segmentation does not produce. Moreover, point cloud segmentation produces poor results in very noisy areas where interpretation and the interpreter's knowledge are essential to extract information.
2. The closest object retrieval approach allows the exploration of different geobodies by providing any kind of object as an example. However, a condition is that the geobody is constituted of a single connected body; broken features in seismic amplitude response, such as salt diapirs or gas chimneys, would not be accurately delineated by the proposed methodology. Seismic interpreters characterise and interpret these broken geobodies from section views; therefore, image processing methods are probably best suited to detect these geobodies.

On seismic interpretation quality, in chapter 3, we discuss the quality of the seismic segmentation and propose metrics to quantify this quality, but it is important in these analyses to keep in mind the non-uniqueness of the seismic interpretation. The non-uniqueness may be due primarily to the uncertainty inherent in the seismic interpretation but also to the choice of segmentation of the interpreter. For example, we discussed in chapter 3 the importance of the notion of hierarchy level for the segmentation of an object: an interpreter may want to interpret a structure or substructure, and each segmentation level leads to different interpretations. HDBSCAN or other hierarchical algorithms could address in part this issue. Therefore, it is not possible to define an absolute quality of interpretation, but these metrics and tools allow us to aim for a good local maximum in terms of quality.

Finally, immersive visualisation is a hot topic with promises to improve the interpretation and understanding of three-dimensional data (figure 6.2), and point cloud, given its sparse nature, is an ideal data format for immersive visualisation. Thus, beyond the detection of geobodies, a point cloud approach could facilitate the visualisation and analysis of seismic data in immersive visualisation.

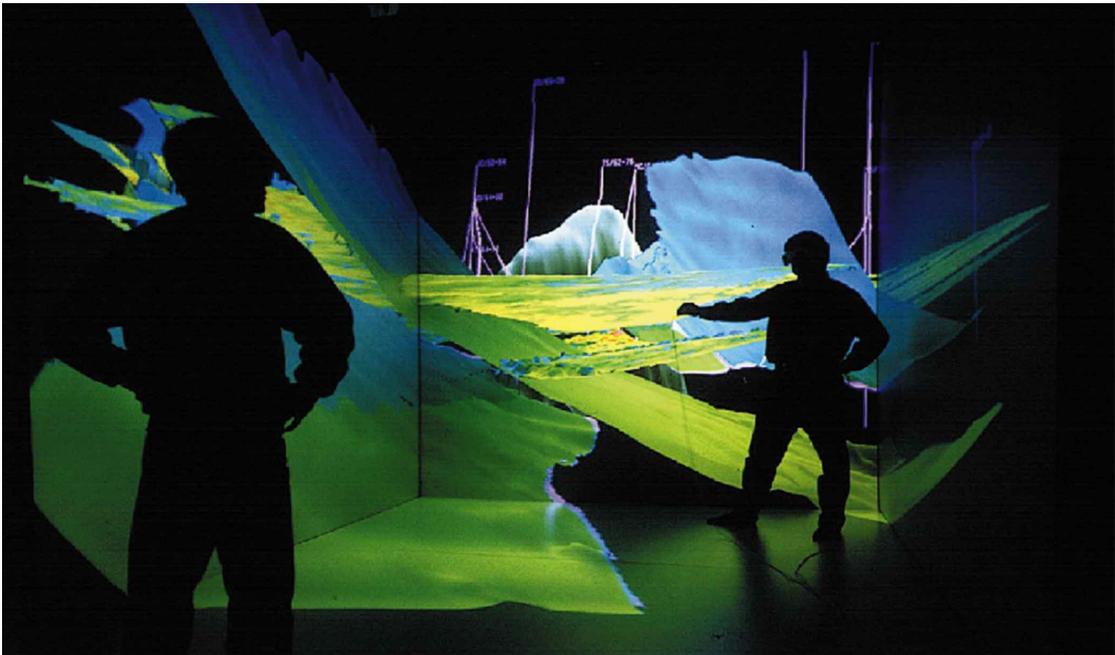


FIGURE 6.2: Visualisation and interpretation of geoscience data in the Immersive Visualisation Environment at ARCO. The three walls and floor are all projection surfaces, and the interpreters are immersed in the data. Figure from Dorn [1998].

6.3 Future work suggestions

Several options exist to expand on the research addressed in this thesis:

1. The first interesting idea would be to integrate and compare different seismic post-processing methods for denoising the seismic and measuring the impact on the seismic segmentation. For instance, Wu [2017] proposed a fault- and stratigraphy-oriented smoothing to enhance the lateral discontinuities corresponding to faults and stratigraphic features in the seismic data and showed that it helps identify subtle features and channels on seismic images. Such smoothing of the seismic

could improve the seismic event continuity and definition and therefore improve the subsequent segmentation, which relies on event continuity.

2. Another potential improvement to the seismic segmentation could be implementing a cut-off value for seismic attributes that evolve with depth. The amplitude response of the seismic reflection data varies with factors such as depth, rock compaction, porosity and lithological composition. Therefore, applying cut-off values that evolve with the local seismic attribute average could improve the segmentation of areas with weak responses.
3. In Chapter 4, we discussed the possibility of characterising unconformities or discordant objects by comparison to the mean local dip of surrounding objects, for instance, but that idea has not been explored in this thesis and would merit further investigation.
4. Another avenue to explore is incorporating notions from other concepts regularly used to characterise geobodies not explored in this thesis. For instance, concepts from spectral decomposition and sequence stratigraphy would add value if embedded in the interpretation. Spectral decomposition allows to better image of geologic boundaries and thin beds that are not fully resolved with seismic amplitude peaks and through and reveals the geometrical details of the seismic object. Thus, working from the spectral decomposition cubes to segment objects would improve the delineation of the geobodies. However, to do this, you have to solve the challenge of working with several cubes -usually three cubes from the spectral decomposition-. In addition, concepts from sequence stratigraphy [Mitchum et al., 1977] could also be included and highlight stratigraphic packages and sequences where it is most likely to encounter certain geobodies, which would reduce the number of possible target objects and increase the confidence in the detection.
5. We could also incorporate data of different natures, such as wireline logs or core data, to get more insights and information to lead the interpretation of the seismic data. For instance, interpretations from well data could be used to label some seismic features, and then we could investigate some semi-supervised machine

learning approaches to perform seismic interpretation from a few labelled samples [Delalleau et al., 2005, Grandvalet and Bengio, 2004].

6. Finally, an interesting extension of this work would be to integrate the geobody extraction in a seismic inversion framework to improve the geobody morphology and predict porosity, lithology and fluid change [Spikes et al., 2007] or even estimate the uncertainties via Stochastic inversions (e.g., Pereira et al. [2020]).

Appendix A

Code available

The code is available for 3D seismic segmentation and geobody detection:

<https://github.com/GeoDataScienceUQ/pyseismic>

A demo is applied to the project F3 Demo 2020 [dGB Earth Sciences, 2020] and presented in a Jupyter notebook.

F3 is a block in the Dutch sector of the North Sea, covered by a 3D seismic survey. The seismic displays a number of interesting features with a large fluvio-deltaic system [Overeem et al., 2001, Schroot and Schüttenhelm, 2003]. Many channels can be observed and detected with our proposed geobody extraction method.

Bibliography

- Aldoma, A., Marton, Z. C., Tombari, F., Wohlking, W., Potthast, C., Zeisl, B., Rusu, R., Gedikli, S., and Vincze, M. (2012). Tutorial: Point cloud library: Three-dimensional object recognition and 6 DOF pose estimation. *IEEE Robotics and Automation Magazine*, 19(3):80–91.
- Alkhalifah, T., Wang, H., and Ovcharenko, O. (2021). MLReal: Bridging the gap between training on synthetic data and real data applications in machine learning. *82nd EAGE Conference and Exhibition 2021*, 7:5478–5482.
- Alsenan, S., Al-Turaiki, I., and Hafez, A. (2020). Autoencoder-based Dimensionality Reduction for QSAR Modeling. *ICCAIS 2020 - 3rd International Conference on Computer Applications and Information Security*.
- Bach, F. R. (2008). Graph kernels between point clouds. *Proceedings of the 25th International Conference on Machine Learning*, pages 25–32.
- Bacon, M., Simm, R., and Redshaw, T. (2007). *3-D Seismic Interpretation*. Cambridge University Press.
- Badley, M. E. (1985). *Practical seismic interpretation*. IHRDC Press, Boston, MA, Pittsburgh, PA, and Morgantown, WV.
- Bahorich, M. and Farmer, S. (1995). 3-D seismic discontinuity for faults and stratigraphic features: The coherence cube. *The Leading Edge*, 14(10):1053–1058.
- Bassier, M., Vergauwen, M., and Poux, F. (2020). Point Cloud vs. Mesh Features for Building Interior Classification. *Remote Sensing 2020, Vol. 12, Page 2224*, 12(14):2224.

- Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517.
- Bentley, M. and Smith, S. (2008). Scenario-based reservoir modelling: The need for more determinism and less anchoring. *Geological Society Special Publication*, 309:145–159.
- Birnie, C. (2021). GEOGRAPHI: AN INTERACTIVE GRAPH DATABASE OF OPENLY AVAILABLE SEISMIC DATASETS. *82nd EAGE Conference and Exhibition 2021*, 7(1):5253–5257.
- Bokulich, A. (2013). Explanatory Models Versus Predictive Models: Reduced Complexity Modeling in Geomorphology. *EPSA11 Perspectives and Foundational Problems in Philosophy of Science*, pages 115–128.
- Bond, C. E., Gibbs, A. D., Shipton, Z. K., and Jones, S. (2007). What do you think this is? "Conceptual uncertainty" In geoscience interpretation. *GSA Today*, 17(11):4–10.
- Borgos, H. G., Skov, T., Randen, T., and Sonneland, L. (2003). Automated geometry extraction from 3D seismic data. *SEG Technical Program Expanded Abstracts*, 22(1):1541–1544.
- Borgos, H. G., Skov, T., and Sønneland, L. (2005). Automated Structural Interpretation Through Classification of Seismic Horizons. In *Mathematics in Industry*, volume volume 7, pages 89–106. Springer, Berlin, Heidelberg.
- Botter, C., Cardozo, N., Hardy, S., Lecomte, I., Paton, G., and Escalona, A. (2016). Seismic characterisation of fault damage in 3D using mechanical and seismic modelling. *Marine and Petroleum Geology*, 77:973–990.
- Bradski, G. and Kaehler, A. (2008). *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly Media, Inc.
- Bugge, A. J., Erik Lie, J., Evensen, A. K., Faleide, J. I., and Clark, S. (2019). Automatic extraction of dislocated horizons from 3D seismic data using nonlocal trace matching. *Geophysics*, 84(6):IM77–IM86.

- Bugge, A. J., Lie, J. E., and Clark, S. (2018). Automatic facies classification and horizon tracking in 3D seismic data. *1st EAGE/PESGB Workshop on Machine Learning*, 2018(1):1–3.
- Bunt, R. J. W. (2015). The use of seismic attributes for fan and reservoir definition in the Sea Lion Field, North Falkland Basin. *Petroleum Geoscience*, 21(2-3):137–149.
- Burgess, P. M., Masiero, I., Toby, S. C., and Duller, R. A. (2019). A Big Fan of Signals? Exploring Autogenic and Allogenic Process and Product In a Numerical Stratigraphic Forward Model of Submarine-Fan Development. *Journal of Sedimentary Research*, 89(1):1–12.
- Campello, R. J., Moulavi, D., and Sander, J. (2013). Density-based clustering based on hierarchical density estimates. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 7819 LNAI, pages 160–172. Springer, Berlin, Heidelberg.
- Canny, J. (1986). A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698.
- Carvalho, D. V., Pereira, E. M., and Cardoso, J. S. (2019). Machine Learning Interpretability: A Survey on Methods and Metrics. *Electronics 2019*, 8(8):832.
- Chang, A. X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., and Yu, F. (2015). ShapeNet: An Information-Rich 3D Model Repository. *arXiv preprint*, arXiv:1512.03012.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal Of Artificial Intelligence Research*, 16:321–357.
- Chen, H. C. and Wang, S. J. (2004). The use of visible color difference in the quantitative evaluation of color image segmentation. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 3.
- Chen, Q. and Sidney, S. (1997). Seismic attribute technology for reservoir forecasting and monitoring. *The Leading Edge*, 16(5):445.

- Chen, Y. S. and Lin, C. Y. (2019). Virtual object replacement based on real environments: Potential application in augmented reality systems. *Applied Sciences (Switzerland)*, 9(9):1797.
- Chevitarese, D., Szwarcman, D., Mozart, R., and Vital Brazil, E. (2018). Seismic Facies Segmentation Using Deep Learning. In *AAPG Annual and Exhibition*. AAPG Annual and Exhibition.
- Chopra, S. and Alexeev, V. (2006). Applications of texture attribute analysis to 3D seismic data. *The Leading Edge*, 25(8):934–940.
- Chopra, S. and Marfurt, K. J. (2005). Seismic attributes - A historical perspective. *Geophysics*, 70(5).
- Cicek, O., Abdulkadir, A., Lienkamp, S. S., Brox, T., and Ronneberger, O. (2016). 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9901 LNCS:424–432.
- Civitarese, D., Szwarcman, D., Brazil, E. V., and Zadrozny, B. (2019). Semantic Segmentation of Seismic Images. *arXiv preprint*, arXiv:1905.04307.
- Clausi, D. A. and Jernigan, M. (1998). A fast method to determine co-occurrence texture features. *IEEE Transactions on Geoscience and Remote Sensing*, 36(1):298–300.
- Clausi, D. A. and Zhao, Y. (2003). Grey level co-occurrence integrated algorithm (GLCIA): a superior computational method to rapidly determine co-occurrence probability texture features. *Computers & Geosciences*, 29(7):837–850.
- Corcoran, P., Winstanley, A., and Mooney, P. (2010). Segmentation performance evaluation for object-based remotely sensed image analysis. *International Journal of Remote Sensing*, 31(3):617–645.
- Corlay, Q., Demyanov, V., McCarthy, D., and Arnold, D. (2020). Turbidite Fan Interpretation in 3D Seismic Data by Point Cloud Segmentation Using Machine Learning. In *EAGE 2020 Annual Conference & Exhibition Online*, volume 2020, pages 1–5. European Association of Geoscientists & Engineers.

- De Bruin, G. and Bouanga, E. G. (2007). Time attributes of stratigraphic surfaces, Analyzed in the structural and wheeler transformed domain. *69th European Association of Geoscientists and Engineers Conference and Exhibition 2007: Securing The Future. Incorporating SPE EUROPEC 2007*, 2:1138–1142.
- De Bruin, G., Ligtenberg, H., Hemstra, N., and Tingdahl, K. (2006). Synchronized sequence stratigraphic interpretation in the structural and chrono-stratigraphic (Wheeler transformed) domain. *Research Workshop 2006: From Seismic Interpretation to Stratigraphic and Basin Modelling - Present and Future*.
- De Groot, P., Huck, A., De Bruin, G., Hemstra, N., and Bedford, J. (2010). The horizon cube A step change in seismic interpretation! *The Leading Edge*, 29(9):1048–1055.
- de Groot, P., Huck, A., de Bruin, G., Hemstra, N., and Bedford, J. (2010). The horizon cube: A step change in seismic interpretation! *The Leading Edge*, 29(9):1048–1055.
- de Groot, P., Pelissier, M., and van Hout, M. (2021). Seismic classification: A thalweg tracking/machine learning approach. *First Break*, 39(3):59–64.
- De Groot, P., Qayyum, F., Liu, Y., and Hemstra, N. (2016). New methods for slicing and dicing seismic volumes. *First Break*, 34(3).
- Delalleau, O., Bengio, Y., and Le Roux, N. (2005). Efficient Non-Parametric Function Induction in Semi-Supervised Learning. *International Workshop on Artificial Intelligence and Statistics*, pages 96–103.
- dGB Earth Sciences (2020). TerraNubis - Data Info of F3 Demo 2020.
- Di, H. (2018). Developing a seismic pattern interpretation network (SpiNet) for automated seismic interpretation. *arXiv preprint*, arXiv:1810.08517.
- Di, H., Wang, Z., and AlRegib, G. (2018). Deep convolutional neural networks for seismic salt-body delineation. *AAPG 2018 Annual Convention & Exhibition*.
- Dodd, T. J. H., McCarthy, D. J., Richards, P. C., and Kane, I. (2019). A depositional model for deep-lacustrine, partially confined, turbidite fans: Early Cretaceous, North Falkland Basin. *Sedimentology*, 66(1):53–80.

- Dorn, G. A. (1998). Modern 3-D seismic interpretation. *The Leading Edge*, 17(9).
- Dramsch, J. S. and Lühje, M. (2018). Deep-learning seismic facies on state-of-the-art CNN architectures. *SEG Technical Program Expanded Abstracts*, pages 2036–2040.
- Du Toit, A. (1920). The Karoo dolerites of South Africa: A study of hypabyssal injection: Geological Society of South Africa Transactions. *Transactions of the Geological Society of South Africa*, 23:1–42.
- Duarte, A. F. (2019). The impact of the Early Cretaceous igneous intrusion in the Falkland Plateau Basin. Technical report, Lisbon University, Lisbon.
- Eitzenberger, A. (2012). *Wave propagation in rock and the influence of discontinuities*. PhD thesis, Lulea tekniska universitet. Institutionen for samhallsbyggnad och naturresurser.
- Elbaz, G., Avraham, T., and Fischer, A. (2017). 3D Point Cloud Registration for Localization Using a Deep Neural Network Auto-Encoder. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4631–4640. CVF.
- Eldholm, O., Skogseid, J., Planke, S., and Gladczenko, T. P. (1995). *Volcanic Margin Concepts*. Springer, Dordrecht.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *kdd*, 96(34):226–231.
- Fernandes, D., Silva, A., Névoa, R., Simões, C., Gonzalez, D., Guevara, M., Novais, P., Monteiro, J., and Melo-Pinto, P. (2021). Point-cloud based 3D object detection and classification methods for self-driving applications: A survey and taxonomy. *Information Fusion*, 68:161–191.
- Fjær, E., Holt, R. M., Horsrud, P., Raaen, A. M., and Risnes, R. (2008). Chapter 5 Elastic wave propagation in rocks. *Developments in Petroleum Science*, 53:175–218.
- Froner, B., Purves, S. J., Lowell, J., and Henderson, J. (2013). Perception of visual information: The role of colour in seismic interpretation. *First Break*, 31(4):29–34.

- Gao, D. (2003). Volume texture extraction for 3D seismic visualization and interpretation. *Geophysics*, 68(4):1294–1302.
- Gao, H., Wu, X., and Liu, G. (2021). ChannelSeg3D: Channel simulation and deep learning for channel interpretation in 3D seismic images. *Geophysics*, 86(4):IM73–IM83.
- Garcia, M., Ercilla, G., Alonso, B., and Estrada, F. (2015). Deep-water turbidite systems: A review of their elements, sedimentary processes and depositional models. Their characteristics on the iberian margins. *Boletín geológico y minero*, 126:189–218.
- Gersztenkorn, A. and Marfurt, K. J. (1999). Eigenstructure-based coherence computations as an aid to 3-D structural and stratigraphic mapping. <https://doi.org/10.1190/1.1444651>, 64(5):1468–1479.
- Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., and Kagal, L. (2019). Explaining explanations: An overview of interpretability of machine learning. *Proceedings - 2018 IEEE 5th International Conference on Data Science and Advanced Analytics, DSAA 2018*, pages 80–89.
- Gionis, A., Indyk, P., and Motwani, R. (1999). Similarity Search in High Dimensions via Hashing. *Vldb*, 99(6):518–529.
- Godefroy, G., Caumon, G., Laurent, G., and Bonneau, F. (2021). Multi-scenario Interpretations From Sparse Fault Evidence Using Graph Theory and Geological Rules. *Journal of Geophysical Research: Solid Earth*, 126(2):e2020JB020022.
- Gogia, R., Gogia, R., Singh, R., De Groot, P., Gupta, H., Srirangarajan, S., Phirani, J., and Ranu, S. (2020). Tracking 3D seismic horizons with a new hybrid tracking algorithm. <http://www.seg.org/interpretation>, 8(4).
- Gong, Y., Lazebnik, S., Gordo, A., and Perronnin, F. (2013). Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929.

- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Networks. *Science Robotics*, 3(January):2672–2680.
- Grandvalet, Y. and Bengio, Y. (2004). Semi-supervised Learning by Entropy Minimization. *Advances in Neural Information Processing Systems*, 17.
- Grilli, E., Poux, F., Remondino, F., and Foundation, B. K. (2021). Unsupervised object-based clustering in support of supervised point-based 3d point cloud classification. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43:471–478.
- Haldorsen, H. H. and Lake, L. W. (1984). A New Approach to Shale Management in Field-Scale Models. *Society of Petroleum Engineers Journal*, 24(04):447–457.
- Hale, D. (2013). Methods to compute fault images, extract fault surfaces, and estimate fault throws from 3D seismic images. *Geophysics*, 78(2).
- Haralick, R. M., Dinstein, I., and Shanmugam, K. (1973). Textural Features for Image Classification. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-3(6):610–621.
- Hart, B. S. (2008). Channel detection in 3-D seismic data using sweetness. *AAPG Bulletin*, 92(6):733–742.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R. (2017). Mask R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2961–2969. CVF.
- Henderson, J., Purves, S. J., Fisher, G., and Leppard, C. (2008). Delineation of geological elements from RGB color blending of seismic attribute volumes. *The Leading Edge*, 27(3):342–350.
- Henderson, J., Purves, S. J., and Leppard, C. (2007). Automated delineation of geological elements from 3D seismic data through analysis of multichannel, volumetric spectral decomposition data. *First Break*, 25(3):87–93.

- Hinton, G. E. and Roweis, S. (2002). Stochastic Neighbor Embedding. *Advances in Neural Information Processing Systems*, 15:857–864.
- Howard, R. E. (1991). METHOD FOR ATTRIBUTE TRACKING IN SEISMIC DATA.
- Hoyes, J. and Cheret, T. (2011). A review of "global" interpretation methods for automated 3D horizon picking. *The Leading Edge*, 30(1):38–47.
- Indyk, P. and Motwani, R. (1998). Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. Association for Computing Machinery (ACM).
- Irakarama, M., Cupillard, P., Caumon, G., Sava, P., and Edwards, J. (2019). Appraising structural interpretations using seismic data — Theoretical elements. *Geophysics*, 84(2):N29–N40.
- Jackson, M. D. and Muggeridge, A. H. (2000). Effect of Discontinuous Shales on Reservoir Performance During Horizontal Waterflooding. *SPE Journal*, 5(04):446–455.
- Jarvis, K. and Saussus, D. (2009). Extracting detailed lithology from seismic data. *ASEG Extended Abstracts*, 2009(1):1.
- Jaynes, E. T. (1968). Prior Probabilities. *IEEE Transactions on Systems Science and Cybernetics*, 4(3):227–241.
- Jin, S., Chen, S., Wei, J., and Li, X. Y. (2017). Automatic seismic event tracking using a dynamic time warping algorithm. *Journal of Geophysics and Engineering*, 14(5):1138–1149.
- Johnson, B. and Xie, Z. (2011). Unsupervised image segmentation evaluation and refinement using a multi-scale approach. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(4):473–483.
- Jones, D. J., McCarthy, D. J., and Dodd, T. J. (2019). Tectonostratigraphy and the petroleum systems in the Northern sector of the North Falkland Basin, South Atlantic. *Marine and Petroleum Geology*, 103:150–162.

- Jones, I. F. and Davison, I. (2014). Seismic imaging in and around salt bodies. *Interpretation*, 2(4):SL1–SL20.
- Kang, H. W., Lee, S. J., Ko, I. K., Kengla, C., Yoo, J. J., and Atala, A. (2016). A 3D bioprinting system to produce human-scale tissue constructs with structural integrity. *Nature Biotechnology*, 34(3):312–319.
- Kanopoulos, N., Vasanthavada, N., and Baker, R. L. (1988). Design of an Image Edge Detection Filter Using the Sobel Operator. *IEEE Journal of Solid-State Circuits*, 23(2):358–367.
- Kington, J. (2015). Semblance, coherence, and other discontinuity attributes. *Leading Edge*, 34(12):1510–1512.
- Labrunye, E. and Carn, C. (2015). Merging chronostratigraphic modeling and global horizon tracking. *Interpretation*, 3(2):SN59–SN67.
- Lacaze, S., Durot, B., Devilliers, A., and Pauget, F. (2017). Comprehensive Seismic Interpretation to Enhance Stratigraphy and Faults. In *15th International Congress of the Brazilian Geophysical Society & EXPOGEF*, pages 1429–1432. Society of Exploration Geophysicists.
- Larsen, E., Purves, S. J., Economou, D., and Alaei, B. (2018). Is machine learning taking productivity in petroleum geoscience on a Moore’s Law trajectory? *First Break*, 36(12):135–141.
- Lawal, M. A., Pecher, I., Bialik, O., Waldmann, N. D., Bialas, J., Koren, Z., and Makovsky, Y. (2022). Multilevel Composition: A new method for revealing complex geological features in three-dimensional seismic reflection data. *Marine and Petroleum Geology*, 146:105938.
- Le Bouteiller, P. and Charléty, J. (2020). Semi-Supervised Multi-Facies Object Retrieval in Seismic Data. *Mathematical Geosciences*, 52(6):817–846.
- Le Bouteiller, P., Lafuerza, S., Charléty, J., Reis, A. T., Granjeon, D., Delprat-Jannaud, F., and Gorini, C. (2019). A new conceptual methodology for interpretation of mass transport processes from seismic data. *Marine and Petroleum Geology*, 103:438–455.

- LeCun, Y. (1989). Generalization and network design strategies. *Connectionism in perspective*, 19:143–155.
- LeCun, Y. and Bengio, Y. (1995). Convolutional Networks for Images, Speech, and Time-Series. *The handbook of brain theory and neural networks*, 3361(10).
- Lecun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- Lev, H. J. (2020). *A Study of 3D Point Cloud Features for Shape Retrieval*. PhD thesis, Université Grenoble Alpes.
- Li, Y., Ku, B., Zhang, S., Ahn, J. K., and Ko, H. (2020). Seismic Data Augmentation Based on Conditional Generative Adversarial Networks. *Sensors*, 20(23):6850.
- Lin, K., Yang, H.-F., Hsiao, J.-H., and Chen, C.-S. (2015). Deep Learning of Binary Hash Codes for Fast Image Retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 27–35. CVPR.
- Lin, Z.-H., Huang, S.-Y. H., and Wang, Y.-C. F. (2020). Convolution in the Cloud: Learning Deformable Kernels in 3D Graph Convolution Networks for Point Cloud Analysis. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1800–1809.
- Liu, H., Wang, R., Shan, S., and Chen, X. (2016). Deep Supervised Hashing for Fast Image Retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2064–2072. CVF.
- Lomask, J., Guitton, A., Fomel, S., Claerbout, J., and Valenciano, A. A. (2006). Flattening without picking. *Geophysics*, 71(4):13–20.
- Long, J., Shelhamer, E., and Darrell, T. (2015). Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440. CVPR.
- Løtveit, Ingrid, F., Gudmundsson, A., and Sydnes, M. (2013). The effect of sill emplacement on petroleum systems. In *Rock fractures in geological processes*. Universitätsdrucke Göttingen.

- Lu, P. (2019). Deep Learning Realm for Geophysics: Seismic Acquisition, Processing, Interpretation, and Inversion. *arxiv preprint*, arXiv:1909.06486.
- Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., and Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1):56–67.
- MacAulay, F. (2015). Sea Lion Field discovery and appraisal: a turning point for the North Falkland Basin. *Petroleum Geoscience*, 21(2-3).
- Mannoor, M. S., Jiang, Z., James, T., Kong, Y. L., Malatesta, K. A., Soboyejo, W. O., Verma, N., Gracias, D. H., and McAlpine, M. C. (2013). 3D printed bionic ears. *Nano Letters*, 13(6):2634–2639.
- Marfurt, K. J. and Kirlin, R. L. (2001). Narrow-band spectral analysis and thin-bed tuning. *Geophysics*, 66(4):1274–1283.
- Marfurt, K. J., Sudhaker, V., Gersztenkorn, A., Crawford, K. D., and Nissen, S. E. (1999). Coherency calculations in the presence of structural dip. *Geophysics*, 64(1):104–111.
- Matheron, G. (1963). Principles of geostatistics. *Economic Geology*, 58(8):1246–1266.
- Maturana, D. and Scherer, S. (2015). VoxNet: A 3D Convolutional Neural Network for real-time object recognition. *IEEE International Conference on Intelligent Robots and Systems*, 2015-December:922–928.
- McInnes, L. (2022). How HDBSCAN Works.
- McInnes, L. and Healy, J. (2017). Accelerated Hierarchical Density Based Clustering. *IEEE International Conference on Data Mining Workshops, ICDMW*, 2017-November:33–42.
- McInnes, L., Healy, J., and Astels, S. (2017). hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11).
- Mian, A. S., Bennamoun, M., and Owens, R. (2006). Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1584–1601.

- Mitchum, R. M. J., Vail, P. R., and Thompson, I. S. (1977). Seismic Stratigraphy and Global Changes of Sea Level: Part 2. The Depositional Sequence as a Basic Unit for Stratigraphic Analysis: Section 2. Application of Seismic Reflection Configuration to Stratigraphic Interpretation. In *AAPG Special Volumes*, volume 165, pages 53–62. AAPG Special Volumes.
- Monier, D., El Rawy, A., and Mahmoud, A. (2021). Delineation of reservoir channels by different seismic attributes and geobody extractions for robust volumetric estimation, Saffron Field, offshore Nile Delta, Egypt. *Leading Edge*, 40(7):484–493.
- Mosser, L., Oliveira, R., and Steventon, M. (2019). Probabilistic seismic interpretation using Bayesian neural networks. *81st EAGE Conference and Exhibition 2019*, 2019(1):1–5.
- Mosser, L., Purves, S., and Naeini, E. Z. (2020). Deep bayesian neural networks for fault identification and uncertainty quantification. *1st EAGE Digitalization Conference and Exhibition*, 2020(1):1–5.
- Müller, M. (2007). Dynamic Time Warping. *Information Retrieval for Music and Motion*, pages 69–84.
- Musgrave, A. W. (1967). SEISMIC REFRACTION PROSPECTING. *SOCIETY OF EXPLORATION GEOPHYSICISTS*.
- Nichols, G. (2009). *Sedimentology and Stratigraphy*. John Wiley & Sons, New York.
- NLOG (2020). TerraNubis - Data Info of F3 Demo 2020.
- Novakovic, D., White, C. D., Corbeanu, R. M., Hammon, W. S., Bhattacharya, J. P., and McMechan, G. A. (2002). Hydraulic Effects of Shales in Fluvial-Deltaic Deposits: Ground-Penetrating Radar, Outcrop Observations, Geostatistics, and Three-Dimensional Flow Modeling for the Ferron Sandstone, Utah. *Mathematical Geology*, 34(7):857–893.
- Omohundro, S. M. (1989). Five Balltree Construction Algorithms. *Berkeley: International Computer Science Institute*, pages 1–22.

- Overeem, I., Weltje, G. J., Bishop-Kay, C., and Kroonenberg, S. B. (2001). The Late Cenozoic Eridanos delta system in the Southern North Sea Basin: A climate signal in sediment supply? *Basin Research*, 13(3):293–312.
- Partyka, G., Gridley, J., and Lopez, J. (1999). Interpretational applications of spectral decomposition in reservoir characterization. *The Leading Edge*, 18(3):353–360.
- Patil, N. K., Malemath, V. S., and Yadahalli, R. M. (2011). Color and texture based identification and classification of food grains using different color models and haralick features. *International Journal on Computer Science and Engineering*, 3(12).
- Paton, G., McArdle, N., Lowell, J., Norton, D., and Purves, S. (2011). Adaptive geobodies: Delineation of complex and heterogeneous stratigraphic features. *SEG Technical Program Expanded Abstracts*, 30:4384–4387.
- Pauget, F., Lacaze, S., and Valding, T. (2009a). A global approach in seismic interpretation based on cost function minimization. In *SEG Technical Program Expanded Abstracts 2009*, pages 2592–2596. Society of Exploration Geophysicists.
- Pauget, F., Lacaze, S., and Valding, T. (2009b). A global approach in seismic interpretation based on cost function minimization. In *SEG Technical Program Expanded Abstracts 2009*, pages 2592–2596. Society of Exploration Geophysicists.
- Paumard, V., Bourget, J., Durot, B., Lacaze, S., Payenberg, T., George, A. D., and Lang, S. (2019). Full-volume 3D seismic interpretation methods: A new step towards high-resolution seismic stratigraphy. *Interpretation*, 7(3):B33–B47.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pereira, P., Calçôa, I., Azevedo, L., Nunes, R., and Soares, A. (2020). Iterative geostatistical seismic inversion incorporating local anisotropies. *Computational Geosciences 2020 24:4*, 24(4):1589–1604.

- Petch, J., Di, S., and Nelson, W. (2022). Opening the Black Box: The Promise and Limitations of Explainable Machine Learning in Cardiology. *Canadian Journal of Cardiology*, 38(2):204–213.
- PetroFAQ Glossary (2022). Geobody definition.
- Pham, N. and Fomel, S. (2020). Uncertainty estimation using bayesian convolutional neural network for automatic channel detection. *SEG Technical Program Expanded Abstracts*, 2020-October:3462–3466.
- Pham, N., Fomel, S., and Dunlap, D. (2019). Automatic channel detection using deep learning. *Interpretation*, 7(3):SE43–SE50.
- Phillips, M. and Fomel, S. (2017). Plane-wave Sobel attribute for discontinuity enhancement in seismic images. *Geophysics*, 82(6):WB63–WB69.
- Planke, S., Alvestad, E., and Eldholm, O. (2012). Seismic characteristics of basaltic extrusive and intrusive rocks. *The Leading Edge*, 18(3):342–348.
- Possee, D., Baines, G., and Liu, Y. (2022). Rapid Assisted Horizon Interpretation Using Multi-Scale Optimization. In *83rd EAGE Annual Conference & Exhibition*, volume 2022, pages 1–5. EAGE Publications BV.
- Poux, F. and Billen, R. (2019). Voxel-based 3D Point Cloud Semantic Segmentation: Unsupervised Geometric and Relationship Featuring vs Deep Learning Methods. *ISPRS International Journal of Geo-Information 2019, Vol. 8, Page 213*, 8(5):213.
- Poux, F., Mattes, C., and Kobbelt, L. (2020). Unsupervised segmentation of indoor 3D point cloud: application to object-based classification. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIV-4(W1-2020):111–118.
- Purves, S. J., Henderson, J., and Leppard, C. (2007). RGB visualisation based delineation of geological elements from volumetric spectral decomposition of 3D seismic data. *69th European Association of Geoscientists and Engineers Conference and Exhibition 2007: Securing The Future. Incorporating SPE EUROPEC 2007*, 5:2812–2816.

- Qayyum, F., De Groot, P., Yasin, J., and Akhter, G. (2012). Building a sequence stratigraphic framework from HorizonCube and well data. *74th European Association of Geoscientists and Engineers Conference and Exhibition 2012 Incorporating SPE EUROPEC 2012: Responsibly Securing Natural Resources*, pages 2205–2209.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2017). PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660.
- Rateau, R., Schofield, N., and Smith, M. (2013). The potential role of igneous intrusions on hydrocarbon migration, West of Shetland. *Petroleum Geoscience*, 19(3):259–272.
- Richards, P., Duncan, I., Phipps, C., Pickering, G., Grzywacz, J., Hoult, R., and Merritt, J. (2006). Exploring for fan and delta sandstones in the offshore Falklands Basins. *Journal of Petroleum Geology*, 29(3):199–214.
- Richards, P. C., Gatliff, R. W., Quinn, M. F., and Fannin, N. G. (1996). PETROLEUM POTENTIAL OF THE FALKLAND ISLANDS OFFSHORE AREA. *Journal of Petroleum Geology*, 19(2):161–182.
- Richards, P. C., Stone, P., Kimbell, G. S., McIntosh, W. C., and Phillips, E. R. (2013). MESOZOIC MAGMATISM IN THE FALKLAND ISLANDS (SOUTH ATLANTIC) AND THEIR OFFSHORE SEDIMENTARY BASINS. *Journal of Petroleum Geology*, 36(1):61–73.
- Richter, R. and Döllner, J. (2010). Out-of-core real-time visualization of massive 3D point clouds. *Proceedings of AFRIGRAPH 2010: 7th International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa*, pages 121–128.
- Richter, R. and Döllner, J. (2014). Concepts and techniques for integration, analysis and visualization of massive 3D point clouds. *Computers, Environment and Urban Systems*, 45:114–124.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695. CVF.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9351:234–241.
- Ruij, J., Caumon, G., and Viseur, S. (2015). Semiautomatic interpretation of 3D sedimentological structures on geologic images: An object-based approach. *Interpretation*, 3(3):SX63–SX74.
- Rutzinger, M., Höfle, B., Hollaus, M., and Pfeifer, N. (2008). Object-Based Point Cloud Analysis of Full-Waveform Airborne Laser Scanning Data for Urban Vegetation Classification. *Sensors 2008, Vol. 8, Pages 4505-4528*, 8(8):4505–4528.
- Sakoe, H. and Chiba, S. (1978). Dynamic Programming Algorithm Optimization for Spoken Word Recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49.
- Schaaf, A. and Bond, E. C. (2019). Quantification of uncertainty in 3-D seismic interpretation: Implications for deterministic and stochastic geomodeling and machine learning. *Solid Earth*, 10(4):1049–1061.
- Schmidt, I., Lacaze, S., and Paton, G. (2013). Spectral decomposition and geomodel interpretation - Combining advanced technologies to create new workflows. *75th European Association of Geoscientists and Engineers Conference and Exhibition 2013 Incorporating SPE EUROPEC 2013: Changing Frontiers*, pages 1309–1313.
- Schnabel, R., Wahl, R., and Klein, R. (2007). Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum*, 26(2):214–226.
- Schofield, N. J., Brown, D. J., Magee, C., and Stevenson, C. T. (2012). Sill morphology and comparison of brittle and non-brittle emplacement mechanisms. *Journal of the Geological Society*, 169(2):127–141.

- Schroot, B. M. and Schüttenhelm, R. T. (2003). Expressions of shallow gas in the Netherlands North Sea. *Netherlands Journal of Geosciences*, 82(1):91–105.
- Schubert, E., Sander, J., Ester, M., Kriegel, H. P., and Xu, X. (2017). DBSCAN Revisited, Revisited: Why and How You Should (Still) Use DBSCAN. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21.
- Senger, K., Millett, J., Planke, S., Ogata, K., Eide, C. H., Festøy, M., Galland, O., and Jerram, D. A. (2017). Effects of igneous intrusions on the petroleum system: a review. *First Break*, 35(6).
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3):379–423.
- Shi, Y., Wu, X., and Fomel, S. (2019). SaltSeg: Automatic 3D salt segmentation using a deep convolutional neural network. <http://www.seg.org/interpretation>, 7(3):SE113–SE122.
- Shirowzhan, S. and Sepasgozar, S. M. (2019). Spatial Analysis Using Temporal Point Clouds in Advanced GIS: Methods for Ground Elevation Extraction in Slant Areas and Building Classifications. *ISPRS International Journal of Geo-Information 2019, Vol. 8, Page 120*, 8(3):120.
- Shwartz-Ziv, R. and Armon, A. (2022). Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90.
- Smith, T. (2017). Geobody interpretation through multiattribute surveys, natural clusters and machine learning. *SEG Technical Program Expanded Abstracts*, pages 2153–2157.
- Son, H. and Kim, C. (2017). Semantic as-built 3D modeling of structural elements of buildings based on local concavity and convexity. *Advanced Engineering Informatics*, 34:114–124.
- Spikes, K., Mukerji, T., Dvorkin, J., and Mavko, G. (2007). Probabilistic seismic inversion based on rock-physics models. *Geophysics*, 72(5):R87–R97.

- Stanca, R. M., McCarthy, D. J., Paton, D. A., Hodgson, D. M., and Mortimer, E. J. (2022). The tectono-stratigraphic architecture of the Falkland Plateau basin; implications for the evolution of the Falkland Islands Microplate. *Gondwana Research*, 105:320–342.
- Stark, T. J. (2004). Relative geologic time (age) volumes - Relating every seismic sample to a geologically reasonable horizon. *Leading Edge*, 23(9):928–932.
- Steffens, B. (2022). *Integrating geological uncertainty and dynamic data into modelling procedures for fractured reservoirs*. PhD thesis, Heriot Watt University, Edinburgh.
- Stephens, T. L., Walker, R. J., Healy, D., Bubeck, A., England, R. W., and McCaffrey, K. J. (2017). Igneous sills record far-field and near-field stress interactions during volcano construction: Isle of Mull, Scotland. *Earth and Planetary Science Letters*, 478:159–174.
- Stone, P., Richards, P. C., Kimbell, G. S., Esser, R. P., and Reeves, D. (2008). Cretaceous dykes discovered in the Falkland Islands: Implications for regional tectonics in the South Atlantic. *Journal of the Geological Society*, 165(1):1–4.
- Strecker, U. and Uden, R. (2002). Data mining of 3D poststack seismic attribute volumes using Kohonen self-organizing maps. *The Leading Edge*, 21(10):1032–1037.
- Suzuki, S. and be, K. A. (1985). Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, 30(1):32–46.
- Sylvester, Z., Durkin, P., and Covault, J. A. (2019). High curvatures drive river meandering. *Geology*, 47(3):263–266.
- Szostak, M. (2020). Automated Land Cover Change Detection and Forest Succession Monitoring Using LiDAR Point Clouds and GIS Analyses. *Geosciences 2020, Vol. 10, Page 321*, 10(8):321.
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., and Liu, C. (2018). A survey on deep transfer learning. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11141 LNCS:270–279.

- Tarsha Kurdi, F. and Awrangjeb, M. (2020). Automatic evaluation and improvement of roof segments for modelling missing details using Lidar data. *International Journal of Remote Sensing*, 41(12):4700–4723.
- Thomson, K. and Hutton, D. (2003). Geometry and growth of sill complexes: insights using 3D seismic from the North Rockall Trough. *Bulletin of Volcanology* 2003 66:4, 66(4):364–375.
- Thomson, K. and Schofield, N. (2008). Lithological and structural controls on the emplacement and morphology of sills in sedimentary basins. *Geological Society Special Publication*, 302:31–44.
- Troya-Galvis, A., Gancarski, P., Passat, N., and Berti-Equille, L. (2015). Unsupervised Quantification of Under-and Over-Segmentation for Object-Based Remote Sensing Image Analysis. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(5):1936–1945.
- Van Der Maaten, L. and Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605.
- Vojodi, H., Fakhari, A., and Eftekhari Moghadam, A. M. (2013). A new evaluation measure for color image segmentation based on genetic programming approach. *Image and Vision Computing*, 31(11):877–886.
- Waldeland, A. and Solberg, A. (2017). Salt Classification Using Deep Learning. In *79th eage conference and exhibition 2017*. EAGE Publications BV.
- Waldeland, A. U., Jensen, A. C., Gelius, L.-J., and Solberg, A. H. S. (2018). Convolutional neural networks for automated seismic interpretation. *The Leading Edge*, 37(7):529–537.
- Walker, R. J., Stephens, T. L., Greenfield, C., Gill, S. P., Healy, D., and Poppe, S. (2021). Segment tip geometry of sheet intrusions, I: Theory and numerical models for the role of tip shape in controlling propagation pathways. *Volcanica*, 4(2):189–201.
- Wan, L., Hurter, S., Bianchi, V., Li, P., Wang, J., and Salles, T. (2022). The roles and seismic expressions of turbidites and mass transport deposits using stratigraphic

- forward modeling and seismic forward modeling. *Journal of Asian Earth Sciences*, 232:105110.
- Wang, Y., Yao, H., and Zhao, S. (2016). Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242.
- Weber, C., Hahmann, S., and Hagen, H. (2010). Sharp feature detection in point clouds. *SMI 2010 - International Conference on Shape Modeling and Applications, Proceedings*, pages 175–186.
- Weiss, Y., Torralba, A., and Fergus, R. (2008). Spectral Hashing. In *Advances in Neural Information Processing Systems*, volume 21.
- Wimmer, M. and Scheiblauer, C. (2006). Instant Points: Fast Rendering of Unprocessed Point Clouds. *PBG@ SIGGRAPH*, pages 129–136.
- Wrona, T., Pan, I., Gawthorpe, R. L., and Fossen, H. (2018). Seismic facies analysis using machine learning. *Geophysics*, 83(5):O83–O95.
- Wu, X. (2017). Directional structure-tensor-based coherence to detect seismic faults and channels. *Geophysics*, 82(2):A13–A17.
- Wu, X., Geng, Z., Shi, Y., Pham, N., Fomel, S., and Caumon, G. (2020). Building realistic structure models to train convolutional neural networks for seismic structural interpretation. *Geophysics*, 85(4):WA27–WA39.
- Wu, X., Liang, L., Shi, Y., and Fomel, S. (2019). FaultSeg3D: Using synthetic data sets to train an end-to-end convolutional neural network for 3D seismic fault segmentation. *Geophysics*, 84(3):IM35–IM45.
- Xiong, W., Ji, X., Ma, Y., Wang, Y., Benhassan, N. M., Ali, M. N., and Luo, Y. (2018). Seismic fault detection with convolutional neural network. *Geophysics*, 83(5):O97–O103.
- Yang, Y., Feng, C., Shen, Y., and Tian, D. (2018). FoldingNet: Point Cloud Auto-Encoder via Deep Grid Deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 206–215. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

- Zeboudj, R. (1988). *Filtrage, seuillage automatique, contraste et contours : du pré-traitement à l'analyse d'image*. PhD thesis, Université de Saint Etienne, France.
- Zhang, H., Fritts, J. E., and Goldman, S. A. (2003). Entropy-based objective evaluation method for image segmentation. In *Storage and Retrieval Methods and Applications for Multimedia 2004*, volume 5307, pages 38–49. SPIE.
- Zhang, J., Zhao, X., Chen, Z., and Lu, Z. (2019). A Review of Deep Learning-Based Semantic Segmentation for Point Cloud. *IEEE Access*, 7:179118–179133.
- Zhao, T. (2018). Seismic facies classification using different deep convolutional neural networks. *SEG Technical Program Expanded Abstracts*, pages 2046–2050.
- Zhou, Q.-Y., Park, J., and Koltun, V. (2018). Open3D: A Modern Library for 3D Data Processing. *arXiv preprint*, arXiv:1801.09847.
- Zhu, H., Long, M., Wang, J., and Cao, Y. (2016). Deep Hashing Network for Efficient Similarity Retrieval. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1):2415–2421.
- Zhu, W., Mousavi, S. M., and Beroza, G. C. (2020). Seismic signal augmentation to improve generalization of deep neural networks. *Advances in Geophysics*, 61:151–177.