Signal-Informed Semi-Supervised 3D Segmentation for Subsurface Analysis

Alexandre Thouvenot*†, Lionel Boillot†, Vincent Gripon*
*IMT Atlantique, LAB-STICC, UMR CNRS 6285, F-29238 Brest, France
†TotalEnergies, OneTech, 64000 Pau, France
Email: * name.surname@imt-atlantique.fr, † name.surname@totalenergies.com

***Abstract*—In this work, we address a well-known challenge in geoscience: the automatic delineation of geological layers within a 3D seismic cube. Traditional approaches, leveraging a limited set of annotated 2D slices, typically fall into two categories: signal-based propagation algorithms or supervised deep learning models. While the former struggle in regions with poor signal quality, the latter fail to generalize to textures that deviate from the annotated reference set. To overcome these limitations, we propose Geo2.5D, a novel method that incorporates an unsupervised regularization loss to better exploit prior knowledge of the signal structure, ensuring improved alignment with the underlying stratigraphic knowledge. Through experiments on a well-established benchmark, we demonstrate the advantages of our approach over conventional semi-supervised methods.**

***Index Terms*—segmentation, 3D image processing, semi-supervised learning, few-shot learning**

I. Introduction

Seismic data, widely used in mining, oil & gas exploration, and civil engineering, provide a representation of subsurface geological structures, typically in the form of 3D images. These images are constructed from seismic reflectors, which are inferred from wave reflections captured by a network of geophones. Estimating these reflectors involves solving an inverse problem to reconstruct the subsurface structure. The resulting data form a 3D seismic cube, which can be visualized as a series of 2D slices, as illustrated in Figure 1. Geoscientists analyze these images to identify various geological features, such as faults, salt domes, channels, or seismic horizons. However, manual interpretation is time-consuming and expensive, motivating extensive research on automating or semi-automating this process [1].

In this work, we focus on the automatic identification of seismic horizons, which correspond to boundaries between geological units—periods of deposition with similar sedimentary characteristics. These horizons are typically identified by changes in acoustic impedance, making them distinguishable in seismic data. Geoscientists analyze these horizons across the seismic cube to better understand the stratigraphic structure of the subsurface. From a computational perspective, identifying seismic horizons can be framed as a 3D image segmentation problem. However, a major challenge in seismic image analysis lies in the high variability of data quality and consistency. Local distortions, discontinuities (e.g. seismic faults), or low signal-to-noise ratios can significantly degrade segmentation performance, motivating the development of specific methods.

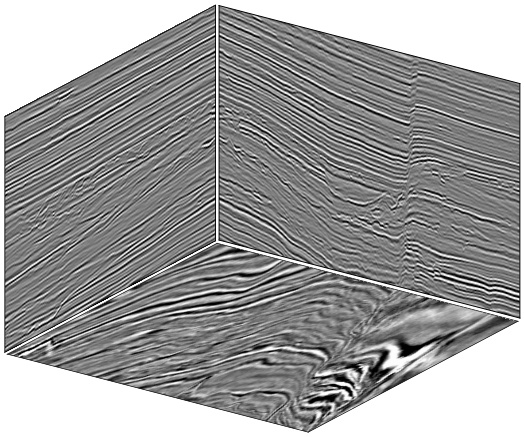

Fig. 1: A depiction of (cropped) frames of the Parihaka 3D seismic cube.

Geoscientists often rely on few-shot learning techniques, where they manually annotate the precise contours of geological layers on a few selected 2D slices of a 3D seismic cube. These annotated slices serve as reference frames to guide the segmentation of the remaining data. Two main approaches exist for this task. The first leverages statistical and signal processing techniques to exploit the limited variation of horizons between adjacent seismic traces. The second leverages Deep Learning (DL) to learn how to associate textures with the corresponding layers. A key limitation of the signal-based strategies is their strong dependence on the quality of the signal. On the other hand, DL-based methods can better handle complex scenarios, such as seismic faults or irregular structures, but their performance remains dependent on the similarity of textures between the reference frames and the target regions. In addition, DL-based strategies often treat 3D segmentation as a series of independent 2D predictions, which can introduce spatial inconsistencies across slices, leading to misaligned geological structures in the final 3D reconstructed prediction cube.

To bridge the gap between these two approaches, we propose a novel semi-supervised training procedure that combines the texture classification capabilities of deep learning models with a signal processing based propagation mechanism to enforce local alignment with seismic stratigraphy. This hybrid

strategy leverages the strengths of both methods, improving segmentation consistency across the 3D seismic cube. We demonstrate that our approach achieves competitive performance on the Parihaka 3D seismic dataset, a well-established benchmark in the field.

II. Related Work

A. Signal-based propagation

Our work draws inspiration from extensive research on local signal alignment. Some of them rely on local signal correlation [2] and provide fast and effective segmentation in areas with clear and continuous geological reflectors. Others are based on seismic local slope [3], [4], where they try to solve differential equations to find surfaces that align with the local wave shift. Even in a local spatial neighborhood, seismic wave patterns may be slightly distorted. Different approaches have considered alignment with slight local deformations such as Dynamic Time Warping (DTW) for horizons propagation [5], or VoxelMorph for 4D seismic alignment [6].

B. Deep Learning Propagation

Several works have explored 3D seismic horizons propagation with DL segmentation models [7]–[9]. These methods are primarily constrained by the availability of labeled data, which is even more costly to obtain in 3D.

Prior works have investigated semi-supervised training for 3D seismic horizon propagation, such as supervised training on few references 2D frames combined with uncertainty reduction in unlabeled areas with contrastive learning [10]. Another possibility is to add physical knowledge as unsupervised regularization as [11]. Moreover, recent work has leveraged seismic slope in a Physic Informed Neural Network (PINN) configuration for seismic data interpolation [12].

Other works have explored the training of foundation models, either with self-supervised pre-training followed by task-specific fine-tuning [13], or with supervised training on both synthetic and real datasets to estimate a seismic attribute, called Relative Geological Time, for horizon propagation [14].

III. Methodology

Let us consider a 3D cube $\mathcal{X}$ of dimensions $N \times H \times W$. We denote $\{\mathbf{X}_1, \dots, \mathbf{X}_N\}$ the N frames of size $H \times W$ along the first dimension. The problem we want to solve can be expressed as follows. We are given a few reference frames $\mathcal{X}_{\text{ref}} = \{\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_L}\}$, typically regularly spaced along the first dimension, with their corresponding segmentation $\mathcal{Y}_{\text{ref}} = \{\mathbf{Y}_{i_1}, \dots, \mathbf{Y}_{i_L}\}$, and we aim to predict the correct segmentation for the remaining frames $\mathcal{X}_{\text{query}} = \mathcal{X} - \mathcal{X}_{\text{ref}}$. Contrary to many other settings in few-shot learning, it is worth pointing out that we have access to all of $\mathcal{X}_{\text{query}}$ for the training phase, what is often referred to as "transductive learning" in the literature [15].

A common baseline to solve this problem consists in training a Deep Learning model f_θ directly on the reference data, what we refer to as the "supervised" method in the following. Namely, the idea is to use a common supervised loss $\mathcal{L}_{\text{sup}}$, such as cross-entropy or DICE for instance, and to solve the optimization problem $\arg\min_\theta \sum_{j=1}^{L} \mathcal{L}_{\text{sup}}(f_\theta(\mathbf{X}_{i_j}), \mathbf{X}_{i_j})$. The predictions are then performed independently of each other on all remaining query frames. As mentioned in the introduction, a main shortcoming of this approach is that it completely disregards the consistency of the predictions along the first dimension, leading to potentially strong misalignment with the underlying seismic stratigraphy.

Our proposed methodology introduces two regularization loss functions, meant to promote a better alignment with the underlying stratigraphy on the query set $\mathcal{X}_{\text{query}}$. The first loss function, denoted $\mathcal{L}_{\text{geo}}$, operates along the third axis, while $\mathcal{L}_{\text{geo}^\perp}$, operates along the first axis. These axes are treated independently because the reference frames being along the first axis disrupt the problem's symmetry. In the coming paragraphs, we delve into more details on these two loss functions.

A. Structure Tensor and Signal Shift

The two proposed regularization loss functions are built on top of an accurate local wave shift estimator. There exist a lot of such estimators in the literature [16]–[19]. We chose the Structure Tensor [17], [18] as it provides a local estimation with a confidence index, well suited to help better tune our loss functions. Denoting i, j and k the indices of our voxels resp. along the first, second and third dimensions, the 2D structure tensor $\mathbf{T}$ is defined as follows:

$$\mathbf{T}_i = \begin{pmatrix} \mathbf{G}_\sigma * \left(\frac{\partial \mathbf{X}_i}{\partial j}\right)^2 & \mathbf{G}_\sigma * \left(\frac{\partial \mathbf{X}_i}{\partial j}\frac{\partial \mathbf{X}_i}{\partial k}\right) \\ \mathbf{G}_\sigma * \left(\frac{\partial \mathbf{X}_i}{\partial j}\frac{\partial \mathbf{X}_i}{\partial k}\right) & \mathbf{G}_\sigma * \left(\frac{\partial \mathbf{X}_i}{\partial k}\right)^2 \end{pmatrix} \tag{1}$$

with $\mathbf{G}_\sigma$ a Gaussian kernel with a standard deviation σ, an hyperparameter adjusted to capture the dominant orientation of the gradient. $\mathbf{T}_i(j,k)$ is a 2×2 symmetric positive semi-definite matrix. We denote its two eigenvalues $\lambda_1^i(j,k) \geq \lambda_2^i(j,k) \geq 0$ with its corresponding orthogonal eigenvectors $\mathbf{u}_i(j,k)$ and $\mathbf{v}_i(j,k)$. We define the coherence C as follow:

$$C_i = \begin{cases} \frac{\lambda_1^i - \lambda_2^i}{\lambda_1^i + \lambda_2^i} \in (0,1) & \text{if } 0 < \lambda_1^i + \lambda_2^i \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

The coherence C gives us a natural signal quality confidence index. If $\lambda_i^2 \ll \lambda_i^1$, $\mathbf{u}_i$ will be well aligned with the gradient direction and we will locally trust the signal shift. To ensure that the predicted horizons align with the seismic stratigraphy, it is necessary to compute the local slope on the prediction as well. Let us denote a prediction $\tilde{\mathbf{Y}} = f_\theta(\mathbf{X})$, which is a $H \times W \times M$ tensor, with M the number of classes. Each element $\tilde{\mathbf{Y}}(j,k)$ is a probability vector indicating the likelihood of belonging to each class. We are interested with the estimated horizons slope, i.e. the boundary slope between two classes. In practice, such an extraction is obtained by applying the $\arg\max$ function on $\tilde{\mathbf{Y}}$ along its third dimension, resulting in a segmentation mask, but this process is non-

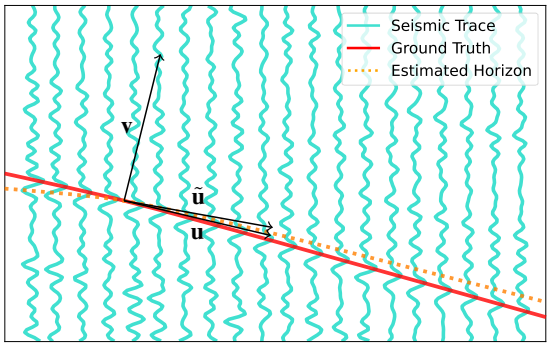

Fig. 2: **Ground-truth and estimated slopes** computed along an horizon.

differentiable. We thus propose to compute an approximation μ of the $\arg\max$ function given by:

$$\mu(\tilde{\mathbf{Y}})(j,k) = (1, \cdots, M)\tilde{\mathbf{Y}}_{j,k} \tag{3}$$

The more closely $\mu(\tilde{\mathbf{Y}})(j,k)$ resembles a one-hot vector, the better it approximates the $\arg\max$ function. As $\mu(\tilde{\mathbf{Y}})$ is a 2D estimated segmentation mask, a closed-form differentiable expression could be derivated from the Structure Tensor to approximate the prediction slope on $f_\theta(\mathbf{X})$. Nonetheless, this method is computationally expensive and simple derivatives are accurate enough. The geological unit boundaries estimated by Equation (3) are quite sharp, resulting in inaccurate slope estimation. Thus, a Gaussian kernel $\mathbf{G}$ with a standard deviation kernel σ' is applied before the computation of the derivative. The estimated horizons local slope vector $\tilde{\mathbf{u}}$ on a prediction $\tilde{\mathbf{Y}}$ is given by:

$$\tilde{\mathbf{u}}_i(j,k) = \left(\frac{\partial(\mu(\tilde{\mathbf{Y}}_i) * \mathbf{G}_{\sigma'})(j,k)}{\partial j}, \frac{\partial(\mu(\tilde{\mathbf{Y}}_i) * \mathbf{G}_{\sigma'})(j,k)}{\partial k} \right) \tag{4}$$

Prediction shifts given by Equation (4) are cheap to compute and differentiable, as both smoothing (Gaussian filter) and derivatives (Sobel filters) can be approximated with 2D convolutions. As shown on Figure 2, the signal structure near one horizon is slowly changing.

B. Stratigraphy alignment

To constrain the overall prediction of the model, we verify that the local seismic slope $\mathbf{u}$ is collinear to the predicted horizon slope $\tilde{\mathbf{u}}$. In a 2D context, this criteria is equivalent to verifying that $\mathbf{v}$ is orthogonal to $\tilde{\mathbf{u}}$, which is more convenient to compute. The regularization loss along the third axis, $\mathcal{L}_{\text{geo}}$ is defined as follows:

$$\mathcal{L}_{\text{geo}}(\mathbf{X}_i, \tilde{\mathbf{Y}}_i) = \frac{1}{H \times W} \sum_{k=1}^{H} \sum_{j=1}^{W} |C_i(j,k)^\alpha \langle \mathbf{v}_i(j,k) | \tilde{\mathbf{u}}_i(j,k) \rangle| \tag{5}$$

with $\alpha \in \mathbb{R}$ a confidence modulation hyperparameter.

Similarly, a regularization loss $\mathcal{L}_{\text{geo}^\perp}$ can be derived where the predictions are still computed on 2D frames along the third axis but the regularization and Structure Tensor are applied along the first one. Let us consider K consecutive 2D frames $\{\mathbf{X}_i, \ldots, \mathbf{X}_{i+K-1}\}$ and their associated predictions $\{\tilde{\mathbf{Y}}_i, \ldots, \tilde{\mathbf{Y}}_{i+K-1}\}$, with K the window size. The geometric loss function $\mathcal{L}_{\text{geo}}$ is computed independently on the truncated orthogonal inputs $\{\mathbf{X}_1^\perp, \ldots, \mathbf{X}_W^\perp\}$ of size $H \times K$ and the predictions $\{\tilde{\mathbf{Y}}_1^\perp, \ldots, \tilde{\mathbf{Y}}_W^\perp\}$ of size $H \times K \times M$. The regularization loss along the first axis, $\mathcal{L}_{\text{geo}^\perp}$, is as follows:

$$\mathcal{L}_{\text{geo}^\perp}\left((\mathbf{X}_j, \tilde{\mathbf{Y}}_j)_{j \in (i, i+K-1)}\right) = \frac{1}{W} \sum_{j=1}^{W} \mathcal{L}_{geo}(\mathbf{X}_j^\perp, \tilde{\mathbf{Y}}_j^\perp) \tag{6}$$

The regularization defined in Equation (6) is computationally expensive to evaluate, as it requires K consecutive predictions. Ideally, we would like $K = W$ as it would provide a global alignment along the entire first direction, anchored by the reference frames, but it would be intractable in practice. To take advantage of the reference frames, while mitigating the computational cost, we work with $5 \leq K \leq 10$ and a custom batch sampling strategy. More specifically, batches are initially sampled close to references frames, and then gradually sampled further away.

C. Training Procedure

The integration of our three objective functions $\mathcal{L}_{\text{sup}}$, $\mathcal{L}_{\text{geo}}$ and $\mathcal{L}_{\text{geo}^\perp}$ require careful consideration, as the latter two are designed to refine the outcomes of the supervised objective. Therefore, we decided to divide the training process into multiple stages, beginning with a initial supervised training phase, followed by two additional regularization phases. Numerous regularization schemes could be designed as long as regularization in different directions does not lead to conflicting objectives. It is often possible for a prediction to be locally aligned with seismic slope in one direction but not in another one. Thus, we chose to start to align the geometry on each query frames with $\mathcal{L}_{\text{geo}}$, and subsequently align query frames between them, with respect to the underlying seismic data, with $\mathcal{L}_{\text{geo}^\perp}$. We call the strategy Geo2.5D, defined by the three following training stages:

1) Pure supervised training with $\mathcal{L}_{\text{sup}}$ on the L images of $\mathcal{X}_{\text{ref}}$. The model f_θ must have converged properly on the few reference frames. The global loss function is $\mathcal{L} = \mathcal{L}_{\text{sup}}$.
2) Supervised with $\mathcal{L}_{\text{sup}}$ on $\mathcal{X}_{\text{ref}}$ and unsupervised training with $\mathcal{L}_{\text{geo}}$ on $\mathcal{X}_{\text{query}}$. The global training loss is $\mathcal{L} = \lambda_{2,\text{sup}}\mathcal{L}_{\text{sup}} + \lambda_{2,\text{geo}}\mathcal{L}_{\text{geo}}$ with $\lambda_{2,\text{sup}} > 0$ and $\lambda_{2,\text{geo}} > 0$. We keep the supervised training objective to prevent an over-regularization that would degrade the overall performance. The training batches are sampled randomly.
3) Supervised with $\mathcal{L}_{\text{sup}}$ on $\mathcal{X}_{\text{ref}}$, unsupervised training with $\mathcal{L}_{\text{geo}}$ on $\mathcal{X}_{\text{query}}$, and $\mathcal{L}_{\text{geo}^\perp}$ on $\mathcal{X}_{\text{query}^\perp}$. The orthogonal loss training set is $\mathcal{X}_{\text{query}^\perp} = \{(\mathbf{X}_i, \cdots, \mathbf{X}_{i+K-1}), i \in (1, N-K+1)\}$, with batches sampled with the strategy mentioned above. The global loss function is $\mathcal{L} = \lambda_{3,\text{sup}}\mathcal{L}_{\text{sup}} + \lambda_{3,\text{geo}}\mathcal{L}_{\text{geo}} + \lambda_{3,\text{geo}^\perp}\mathcal{L}_{\text{geo}^\perp}$ with $\lambda_{3,\text{sup}} > 0$, $\lambda_{3,\text{geo}} > 0$ and $\lambda_{3,\text{geo}^\perp} > 0$.

IV. Results

In this section, we evaluate our method on a classical benchmark in the field, the Parihaka dataset [20]. As our work explores slope based regularization, we used a modified version of it, presented in [11], in which the geological units follow locally the seismic slope. The seismic cube has a size of $N \times H \times W = 777 \times 590 \times 1006$, and the corresponding segmentation is composed of 7 different classes. Our reference dataset is composed of 3 frames $\mathcal{X}_{\text{ref}} = \{\mathbf{X}_{192}, \mathbf{X}_{384}, \mathbf{X}_{576}\}$. Our method, Geo2.5D, alongside its variations tested out in the following sections, are presented with U-Net model [21], albeit the method is architecture-agnostic.

A. Regularization stages ablation

To evaluate each training steps of our method, we propose to compare Geo2.5D with three ablated variations, as reported in Table I. The overall number of epochs was adjusted for each experience to fully converge. We selected two classical segmentation metrics, mIOU and F1, to assess the segmentation results. To evaluate the stability along the first direction, we compute metric h, which is defined as follows:

$$h(n) = \frac{1}{HW} \left| \left\| \mathbb{1}_{\mathbb{Z}^*}(\mathbf{Y}_n - \mathbf{Y}_{n+1}) \right\|^2 - \left\| \mathbb{1}_{\mathbb{Z}^*}(\tilde{\mathbf{Y}}'_n - \tilde{\mathbf{Y}}'_{n+1}) \right\|^2 \right| \tag{7}$$

where $\mathbb{1}_{\mathbb{Z}^*}$ is the element-wise indicator function of $\mathbb{Z}^*$ and $\tilde{\mathbf{Y}}' = \arg\max \tilde{\mathbf{Y}}$ is the predicted segmentation mask. As the seismic data, and the corresponding segmentation mask, have a progressive spatial evolution, we expect the predictions the have a similar variability. Note that the h metric alone does not reflect the quality of segmentation. We report in Table I the performance of the different configurations. First, the two regularization losses, when applied individually, improve the segmentation results and reduce the spatial instabilities along the first direction. Nonetheless, the orthogonal regularization (+8.3% of mIOU compared to the baseline) outperforms the second configuration (+7.0% of mIOU compared to the baseline), presumably as it benefits from reference frames anchor points. Furthermore, combining the regularization losses in Geo2.5D, surpasses all 3 other experiments (+10.1% of mIOU compared to the baseline).

TABLE I: Evaluation of the performances of each regularization losses on the overall results.

$\mathcal{L}_{sup}$	$\mathcal{L}_{geo}$	$\mathcal{L}_{\text{geo}^\perp}$	mIOU ↑	F1 ↑	$\bar{h}$ ($\times 10^{-3}$) ↓
✓			87.1	92.5	9.04
✓	✓		93.2	96.4	2.39
✓		✓	94.3	96.9	2.24
✓	✓	✓	**95.9**	**97.9**	**0.54**

B. Regularization losses ablations

We report in Table II the comparison of Geo2.5D with using Random Sampling (RS) for $\mathcal{X}_{\text{query}^\perp}$. Custom sampling surpasses Random Sampling in segmentation scores, whereas Random Sampling seems to provide a similar spatial stability. We hypothesize that, as Random Sampling does not rely on reference frames as anchors, the local alignment may often be accurate with the seismic slope but misaligned with the reference frames.

TABLE II: Performance on different sampling strategies for $\mathcal{X}_{\text{query}^\perp}$.

	mIOU ↑	F1 ↑	$\bar{h}$ ($\times 10^{-3}$) ↓
Geo2.5D RS	94.9	97.3	**0.50**
Geo2.5D	**95.9**	**97.9**	0.54

C. Benchmark Parihaka

To better evaluate our method, we selected a standard supervised baseline, corresponding to the first setup in Table I, along with two other competitive approaches: a contrastive based method CONSS [10] and a semi-supervised method based on an empirical stratigraphic stacking assumption presented in [11]. We report in Table III the comparison of results on the considered Parihaka dataset. Our method outperforms both semi-supervised ones while providing a better spatial alignment, emphasizing the efficacy of slope-based regularization.

TABLE III: Performance on modified Parihaka dataset. The computational times were measured using a single A100 GPU.

	mIOU ↑	F1 ↑	$\bar{h}$ ($\times 10^{-3}$) ↓	Time (min)
Supervised	87.1	92.5	9.04	3
Stratigraphic [11]	91.6	95.4	3.19	11
CONSS [10]	95.0	97.4	1.76	45
Geo2.5D (ours)	**95.9**	**97.9**	**0.54**	41

We compare in Figure 3 two frames of the predicted 3D cube, with the four selected methods, one of the query dataset and one along the first axis. Although the stratigraphic method removes major layer inconsistencies, both the baseline and the stratigraphic methods demonstrate limited accuracy in predicting horizons in both directions. Our method provides an overall better approximation in both directions, while CONSS shows superior performance in areas with poor slope confidence (bottom horizon of yellow class).

V. Limitations & Future Work

a) Blurry interfaces: depending on the task, geoscientists may look for geological units with interfaces not delineated by a relatively clear reflector, which is often the case in eroded areas. In such conditions, our method tends to underperform as the Structure Tensor and its confidence may not be reliable.

b) Multi-Directional regularization: seismic interpretation often requires examining the 3D cube along multiple directions to fully grasp the stratigraphic structure. A natural enhancement of our method would be to leverage the stratigraphical alignment on multiple 2D plane, similar to [22] but with regularization losses.

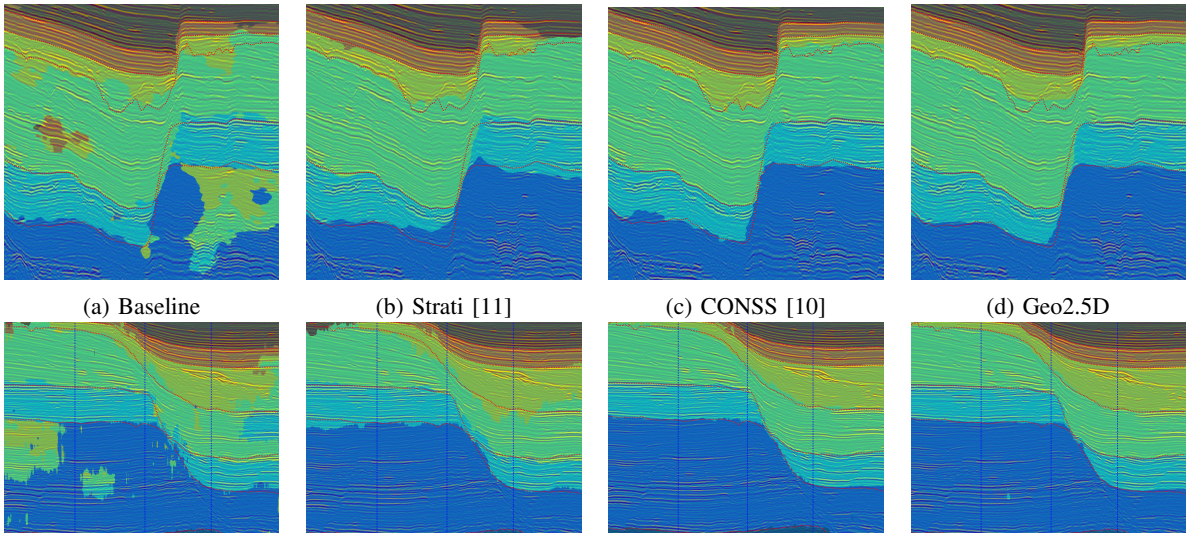

(a) Baseline (b) Strati [11] (c) CONSS [10] (d) Geo2.5D

(e) Baseline (f) Strati [11] (g) CONSS [10] (h) Geo2.5D

Fig. 3: **Benchmark predictions** with the 4 different cases. The first row (a-d) shows the results on query frames. The second row (e-h) presents the results on the orthogonal prediction plane. Red dotted lines are the ground truth horizons and the blue dotted lines are the positions of references frames along the orthogonal direction.

VI. Conclusion

In this paper, we address the challenge of the 3D segmentation of seismic cubes, using a few 2D frames as references. We proposed a method leveraging both signal-based properties of seismic data, specifically the Structure Tensor, and a DL model to provide predictions in a 2.5D manner. Our experiments suggest that our combined strategy can achieve better performance on competitive benchmarks compared to classical alternatives based on supervised or semi-supervised learning.

VII. Acknowledgements

Authors would like to thank New Zealand Petroleum and Minerals for the public Parihaka seismic dataset.

References

[1] R. Chen, Z. Zhang, and J. Ma, "Seismic fault sam: Adapting sam with lightweight modules and 2.5d strategy for fault detection," *arXiv preprint arXiv:2407.14121*, 2024.

[2] J. Gallon, S. Guillon, B. Jobard, H. Barucq, and N. Keskes, "Slimming brick cache strategies for seismic horizon propagation algorithms," *Eurographics/IEEE VGTC on Volume Graphics*, 05 2010.

[3] G. Zinck, M. Donias, J. Daniel, and O. Lavialle, "Local dip transformation for fast seismic horizon reconstruction," in *2013 IEEE International Conference on Image Processing*, 2013, pp. 2368–2372.

[4] S. Guillon, N. Keskes, J. Gallon, and M. Donias, "Geotime: A 3d automatic tool for chronostratigraphic seismic interpretation and filtering," *The Leading Edge*, vol. 32, no. 2, pp. 154–159, 2013.

[5] M. Su, F. Qian, S. Cui, C. Yuan, and X. Cui, "Research on a 3d seismic horizon automatic-tracking method based on corrugated global diffusion," *Applied Sciences*, vol. 13, no. 10, 2023.

[6] J. S. Dramsch, A. N. Christensen, C. MacBeth, and M. Lüthje, "Deep unsupervised 4-d seismic 3-d time-shift estimation with convolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2022.

[7] R. M. Silva, L. Baroni, R. S. Ferreira, D. Civitarese, D. Szwarcman, and E. V. Brazil, "Netherlands dataset: A new public dataset for machine learning in seismic interpretation," *arXiv preprint arXiv:1904.00770*, 2019.

[8] Y. Alaudah, P. Michalowicz, M. Alfarraj, and G. AlRegib, "A machine learning benchmark for facies classification," *arXiv preprint arXiv:1901.07659*, 2019.

[9] V. Tschannen, M. Delescluse, N. Ettrich, and J. Keuper, "Extracting horizon surfaces from 3D seismic data using deep learning," *Geophysics*, vol. 85, no. 3, pp. N17–N26, 2020.

[10] K. Li, W. Liu, Y. Dou, Z. Xu, H. Duan, and R. Jing, "Conss: Contrastive learning method for semisupervised seismic facies classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 7838–7849, 2023.

[11] L. Boillot, A. Thouvenot, J. K. de Chizelle, and S. Guillon, *Stratigraphic constraint for deep learning image segmentation of geological units*, 2023, pp. 763–767.

[12] F. Brandolin, M. Ravasi, and T. Alkhalifah, "Pinnslope: Seismic data interpolation and local slope estimation with physics informed neural networks," *GEOPHYSICS*, vol. 89, no. 4, pp. V331–V345, 2024.

[13] H. Sheng, X. Wu, X. Si, J. Li, S. Zhang, and X. Duan, "Seismic foundation model: A next generation deep-learning model in geophysics," *GEOPHYSICS*, vol. 90, no. 2, pp. IM59–IM79, 2025.

[14] Z. Bi, X. Wu, Z. Geng, and H. Li, "Deep relative geologic time: A deep learning method for simultaneously interpreting 3-d seismic horizons and faults," *Journal of Geophysical Research: Solid Earth*, vol. 126, no. 9, 2021.

[15] O. Chapelle, B. Schölkopf, and A. Zien, Eds., *Semi-supervised learning*, ser. Adaptive computation and machine learning series. Cambridge, Mass. [u.a]: MIT Press, 2010.

[16] D. Hale, *Fast local cross-correlations of images*, 2006, pp. 3160–3164.

[17] G. C. Fehmers and C. F. W. Höcker, "Fast structural interpretation with structure-oriented filtering," *GEOPHYSICS*, vol. 68, no. 4, pp. 1286–1293, 2003.

[18] W. Förstner and E. Gülch, "A fast operator for detection and precise location of distinct points, corners and centres of circular features," in *Proc. ISPRS intercommission conference on fast processing of photogrammetric data*, vol. 6. Interlaken, 1987, pp. 281–305.

[19] Y. Lou, S. Li, N. Liu, and R. Liu, "Seismic volumetric dip estimation via a supervised deep learning model by integrating realistic synthetic data sets," *Journal of Petroleum Science and Engineering*, vol. 218, p. 111021, 2022.

[20] "Seismic facies identification challenge," accessed March 2025. [Online]. Available: https://www.aicrowd.com/challenges/seismic-facies-identification-challenge#links

[21] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *arXiv preprint arXiv:1505.04597*, 2015.

[22] M. Perslev, E. B. Dam, A. Pai, and C. Igel, *One Network to Segment Them All: A General, Lightweight System for Accurate 3D Medical Image Segmentation*. Springer International Publishing, 2019, p. 30–38.