applied sciences

MDPI

Article

Redundancy-Aware Analysis of Functional Complementarity in Seismic Attributes for Deep Facies Segmentation

Roberto Carlos Moreno-Hernández [1,*], Juan A. Moreno-Hernández [2,*], Margarita De la Portilla-Reynoso [3,*], Claudia del C. Gutiérrez-Torres [2], Juan G. Barbosa-Saldaña [2], Didier Samayoa [2] and José A. Jiménez-Bernal [2]

[1] División de Geociencias Aplicadas, Instituto Potosino de Investigación Científica y Tecnológica A.C., San Luis Potosí 78216, Mexico
[2] Escuela Superior de Ingeniería Mecánica y Eléctrica, Instituto Politécnico Nacional, Mexico City 07738, Mexico; cgutierrezt@ipn.mx (C.d.C.G.-T.); jjimenezb@ipn.mx (J.A.J.-B.)
[3] Escuela Superior en Economía, Instituto Politécnico Nacional, Mexico City 11350, Mexico
* Correspondence: roberto.moreno@ipicyt.edu.mx (R.C.M.-H.); jmorenoh1800@alumno.ipn.mx (J.A.M.-H.); mdelaportillar2100@alumno.ipn.mx (M.D.l.P.-R.)

Abstract

Seismic attribute selection remains a critical yet often heuristic component in deep learning-based segmentation workflows. In this work, we propose a redundancy-aware framework to systematically analyse the contribution of seismic attributes by combining input-space statistics, representational similarity (CKA), and error-based evaluation. Our results suggest that statistical redundancy in the input space does not directly translate to functional redundancy within the network. In particular, attributes such as amplitude and instantaneous phase may exhibit high similarity in the input space while producing distinct error patterns and meaningful performance gains. We further observe that complementary attributes do not necessarily yield additive improvements. While some combinations introduce conflicting interactions that limit global performance, others provide stable and consistent improvements across classes. Notably, the combination of amplitude, phase, and local variance forms a minimal informative subset that improves segmentation performance in a balanced manner, particularly in challenging facies. These findings suggest that attribute selection should be guided by functional complementarity and interaction stability rather than by input diversity alone. The proposed framework provides a principled approach for identifying effective attribute subsets, contributing to more efficient and interpretable seismic segmentation workflows.

Keywords: seismic attributes; facies segmentation; functional redundancy; redundancy-aware framework; Centered Kernel Alignment; U-Net

Academic Editor: Nikolaos Koukouzas

Received: 12 May 2026
Revised: 18 June 2026
Accepted: 21 June 2026
Published: 23 June 2026

1. Introduction

Seismic interpretation is essential for subsurface characterisation in hydrocarbon exploration, reservoir development, and regional geological studies. Despite major advances in seismic acquisition and imaging, interpreting large three-dimensional surveys remains time-consuming and heavily dependent on interpreter experience and subjective judgement [1–3]. As seismic datasets continue to grow in complexity and size, there is increasing demand for quantitative workflows that improve reproducibility and consistency.

Seismic attributes have long been used to enhance interpretation by emphasising different physical aspects of the signal. Attributes derived from amplitude, phase, frequency, structural continuity, and local texture can reveal geological features such as stratigraphic

boundaries, channel systems, faults, and facies transitions that may be subtle or invisible in raw amplitudes [4–6]. However, the proliferation of available attributes has introduced a methodological challenge: in practice, attribute selection remains largely heuristic. Interpreters often combine commonly used attributes without systematically evaluating their informational contribution or redundancy. Consequently, determining which attributes provide genuinely complementary information remains an open question.

Recent advances in deep learning offer new possibilities for automated seismic interpretation, particularly through convolutional neural networks for semantic segmentation [7]. Architectures such as U-Net have achieved strong performance in seismic facies segmentation by learning spatially coherent representations directly from seismic images [8,9]. Beyond improving accuracy, these models provide a controlled analytical framework for evaluating how different seismic attributes influence model behaviour. Critically, because deep neural networks learn nonlinear feature representations, attributes that appear statistically similar in the input space may still provide functionally complementary information when processed through deep learning networks.

Despite the growing adoption of deep learning in seismic interpretation, systematic evaluations of attribute contribution and redundancy within deep segmentation frameworks remain limited. In many existing studies, attributes are simply concatenated as additional input channels without quantifying their marginal contribution or redundancy relative to other attributes. As a result, attribute selection in deep-learning-based workflows is often guided by convention rather than by quantitative evidence. This gap is particularly problematic because traditional redundancy measures (e.g., correlation, mutual information) operate in the input space and cannot account for the nonlinear transformations learned by deep networks.

This study addresses this gap by introducing a redundancy-aware quantitative framework for evaluating seismic attributes in deep facies segmentation. The work is built on the central hypothesis that statistical redundancy between seismic attributes does not necessarily imply functional redundancy within nonlinear segmentation models. Under this premise, attributes that appear redundant in input space may still provide complementary information when processed through deep neural networks.

Using a controlled experimental protocol based on the publicly available F3 Netherlands seismic benchmark [9,10], we analyse a set of widely used attributes spanning amplitude-based, phase-based, structural, and texture-related categories. A fixed segmentation architecture and consistent training configuration are maintained across all experiments to isolate the effect of attribute inputs and ensure reproducibility.

The main contributions of this work are summarised as follows:

1. A redundancy-aware quantitative framework for evaluating seismic attributes in deep facies segmentation under controlled experimental conditions.
2. A systematic analysis of marginal attribute contribution, measured through controlled ablation experiments relative to an amplitude-only baseline.
3. A comparative evaluation of statistical and functional redundancy among commonly used seismic attributes, highlighting cases where statistical similarity does not imply redundant segmentation behaviour.
4. The identification of minimal informative attribute subsets that provide complementary information for facies discrimination while avoiding unnecessary redundancy.

By clarifying which seismic attributes provide genuinely complementary information in nonlinear segmentation models, this work aims to provide practical guidance for attribute selection in modern seismic interpretation workflows and to support more transparent and reproducible applications of deep learning in applied geophysics.

2. Related Work

Seismic attributes have long been fundamental in post-stack interpretation workflows for enhancing geological information in reflection data. Early work demonstrated that attributes derived from the analytic signal—instantaneous amplitude, phase, and frequency—highlight stratigraphic and structural features not immediately apparent in raw amplitudes [4]. Subsequent expansions included reflector continuity, curvature, and local texture, providing interpreters with broad tools for identifying geological structures and depositional patterns [5,6]. Nevertheless, attribute selection in practice remains largely guided by empirical experience rather than systematic quantitative evaluation.

In parallel, feature selection and redundancy analysis are well established in machine learning. Techniques based on correlation, mutual information, and wrapper methods are commonly used to identify informative, non-redundant features. In geoscience, several studies have explored attribute ranking and selection strategies to improve machine-learning-based facies classification [11,12]. However, these approaches are primarily developed for conventional machine learning models and do not account for the nonlinear feature transformations learned by deep neural networks.

Deep learning has been widely adopted for seismic image analysis, particularly semantic segmentation. Convolutional neural networks, especially encoder–decoder architectures such as U-Net, have demonstrated strong performance in seismic facies segmentation by learning spatially coherent representations from seismic sections [8,9]. The availability of public benchmarks, such as the F3 Netherlands survey, has enabled systematic comparisons between models and workflows [10]. Recent reviews highlight rapid progress in deep-learning-based seismic interpretation while emphasising the lack of standardised evaluation protocols and reproducible experimental designs [7,13].

Recent studies have explored architectural improvements for seismic facies segmentation, including attention mechanisms, attention-guided encoder–decoder architectures, and dilated convolutional networks. These approaches improve feature extraction and boundary delineation, leading to enhanced segmentation performance [14–16]. However, these works primarily focus on network design rather than systematically analysing the informational contribution of input attributes.

Despite these advances, attribute selection in deep-learning-based segmentation workflows remains largely heuristic. Existing studies have explored statistical attribute selection using feature-ranking, correlation analysis, mutual information, explainable machine learning, and conventional feature-selection algorithms to improve seismic facies classification [12,17–19]. However, these approaches evaluate attribute relevance and redundancy exclusively in the input feature space and do not account for the nonlinear feature representations learned by deep neural networks. Consequently, in many deep-learning-based segmentation workflows, multiple seismic attributes are concatenated as input channels without explicitly quantifying their marginal contribution or functional redundancy. As a result, it remains unclear whether statistically similar attributes provide redundant or complementary information when processed through deep nonlinear segmentation models.

This limitation motivates the redundancy-aware evaluation framework proposed in this study, which systematically quantifies attribute contribution and redundancy under controlled deep facies segmentation conditions.

3. Dataset and Fixed Experimental Protocol

3.1. F3 Netherlands Seismic Dataset

Experiments are conducted using the publicly available F3 Netherlands seismic survey, a three-dimensional offshore dataset located in the Dutch sector of the North Sea. The survey covers approximately 387 km^2 and is distributed as part of an open seismic

interpretation benchmark widely used in machine-learning-based facies segmentation studies [9,10]. The dataset consists of a 3D post-stack seismic amplitude volume and its corresponding voxel-wise facies labels. Following the standard benchmark protocol, the 3D volume is processed as individual 2D inline and crossline sections, which constitute the input to the U-Net. The dataset includes interpreted stratigraphic horizons that define six facies units: Upper North Sea, Middle North Sea, Lower North Sea, Rijnland/Chalk, Scruff, and Zechstein.

The dataset exhibits strong class imbalance, with Lower North Sea facies dominating most samples, while Scruff and Zechstein occur much less frequently. Such imbalance is typical in stratigraphic interpretation problems and represents an important challenge for automated facies classification.

3.2. Preprocessing and Normalization

Seismic sections are resized to 512×256 pixels to balance spatial context and computational efficiency. Continuous input channels are resampled using bilinear interpolation, while facies labels are resampled using nearest-neighbour interpolation to preserve discrete class identities.

All input channels are standardised using per-channel z-score normalisation. The mean and standard deviation are computed exclusively from the training split and subsequently applied to the validation and test sets to avoid information leakage.

During training, data augmentation is applied to improve model generalisation. The augmentation strategy includes horizontal flips, small rotations, and mild brightness and contrast adjustments applied jointly to seismic inputs and facies masks. Validation and test datasets undergo deterministic preprocessing only.

3.3. Fixed U-Net Backbone Architecture

A U-Net encoder–decoder architecture [8] is used as the segmentation backbone throughout the study. U-Net architectures are widely used for dense prediction tasks because skip connections allow the network to combine multi-scale contextual information with precise spatial localisation.

The network consists of four encoder levels with max-pooling operations for downsampling and corresponding decoder stages using transposed convolutions for upsampling. Skip connections concatenate encoder and decoder feature maps at matching spatial resolutions (Figure 1). The network outputs six class logits corresponding to the facies units defined in the dataset.

Using a fixed architecture ensures that segmentation performance differences observed in this study arise primarily from variations in the input attribute configuration rather than from architectural modifications.

3.4. Model Training Setup

All models are trained under identical optimisation settings to isolate the effect of seismic attributes. The network is trained from scratch using the Adam optimiser [20] with class-weighted cross-entropy to address class imbalance.

Training hyperparameters are summarised in Table 1.

Each configuration is evaluated while keeping the architecture, preprocessing, and training settings fixed, ensuring that performance differences are attributable solely to the informational contribution of the evaluated attributes.

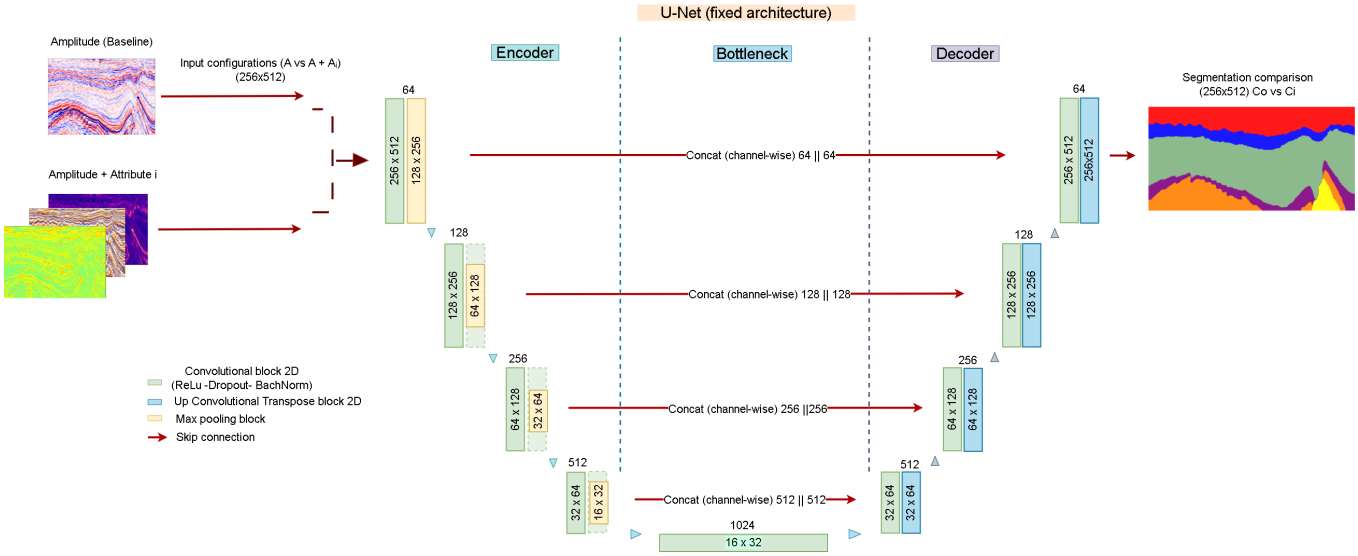

Figure 1. Architecture used as the fixed segmentation backbone. Input configurations include amplitude-only (baseline) and amplitude combined with a single seismic attribute, enabling controlled evaluation of marginal attribute contribution. The encoder–decoder structure incorporates skip connections that concatenate feature maps at matching spatial resolutions. The network outputs a 512×256 facies segmentation map with six classes, where each color represents a different facies class.

Table 1. Training configuration and hyperparameters.

Component	Setting
Dataset	F3 Netherlands benchmark
Input channels	1 (baseline); 2 (amplitude + one attribute)
Attributes	Phase: Instantaneous phase, instantaneous frequency Energy: envelope, RMS amplitude Texture: local variance, entropy Structural: coherence
Backbone	U-Net encoder–decoder
Initial channels	64 (doubling with depth)
Output	1×1 convolution $\rightarrow$ 6 class logits
Loss function	Class-weighted cross-entropy
Class weights	$w_c = 1/(f_c + \epsilon)$ (normalised)
Optimizer	Adam
Batch size	8
Learning rate	10^{-4}
Epoch budget	75
Regularisation	BatchNorm + dropout
Data augmentation	Joint transforms
Seed	42

3.5. Methodology

The supervised segmentation workflow follows the fixed protocol described in Sections 3.1–3.4. Inputs are either amplitude-only or amplitude combined with a single seismic attribute, co-registered on the seismic grid and processed identically across all experiments. The F3 Netherlands survey is partitioned into non-overlapping training (northwest), validation (east), and test (southwest) regions following the official benchmark protocol, thereby preventing spatial leakage between subsets. During training, data augmentation is applied on-the-fly to increase sample diversity while preserving the original class distribution. The five-stage pipeline—(1) input preparation and feature engineering,

(2) training, (3) validation and model selection, (4) testing, and (5) attribute contribution analysis—is illustrated in Figure 2. Because the architecture, preprocessing, and training configuration remain fixed, and all models are evaluated using the final checkpoint after a fixed epoch budget, any performance difference is attributable solely to the informational contribution of the evaluated attributes under identical optimisation conditions.

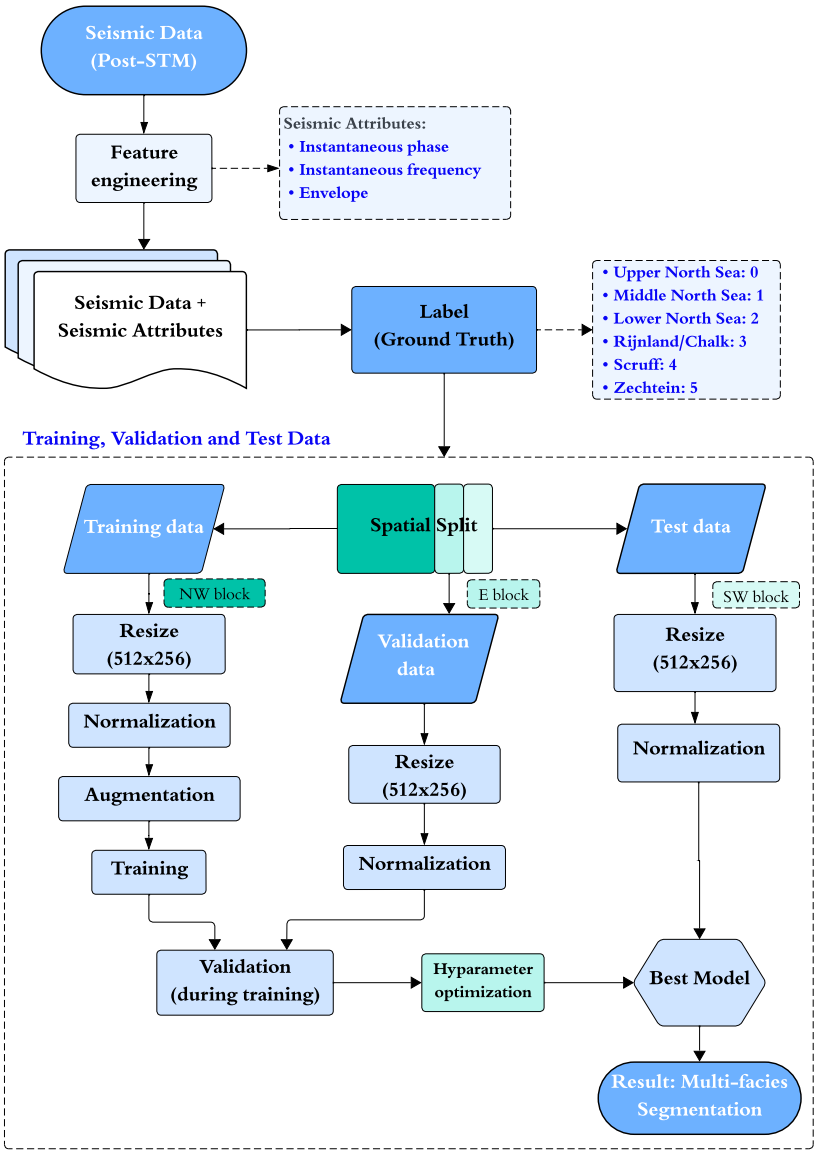

Figure 2. Workflow of the proposed redundancy-aware framework for seismic facies segmentation. Seismic amplitude is used as the baseline input and combined with a single seismic attribute to form two-channel inputs, enabling controlled evaluation of marginal attribute contribution. The F3 Netherlands survey is spatially partitioned into training (NW), validation (E), and test (SW) regions. The fixed U-Net backbone is trained with data augmentation, model selection is performed on the validation set, and final performance is evaluated on the held-out test region.

4. Redundancy-Aware Attribute Evaluation Framework

This section presents the methodological core of the study. A redundancy-aware evaluation framework is proposed to quantify the informational contribution of seismic attributes under controlled deep facies segmentation conditions. The framework integrates input-space redundancy analysis, controlled marginal contribution experiments, functional redundancy assessment, and efficiency-aware subset selection.

4.1. Attribute Categories and Computation

Seven seismic attributes are considered in addition to the amplitude baseline and grouped into physically interpretable categories commonly used in seismic interpretation [4–6]:

- Baseline: seismic amplitude $x(t)$.
- Phase-based: instantaneous phase, instantaneous frequency.
- Energy-based: envelope, Root Mean Square (RMS) amplitude.
- Texture and complexity: local variance (normalised), entropy (computed on the envelope).
- Structural: coherence (semblance-based).

Instantaneous attributes are derived from the analytic signal formulation widely used in time–frequency signal analysis [4,21]:

$$z(t) = x(t) + j\mathcal{H}[x(t)] \tag{1}$$

where $\mathcal{H}[\cdot]$ denotes the Hilbert transform. The envelope and instantaneous phase are defined as:

$$A(t) = |z(t)|, \quad \phi(t) = \arg(z(t)) \tag{2}$$

Instantaneous frequency is computed as:

$$\omega(t) = \frac{d}{dt}\,\mathrm{unwrap}(\phi(t)) \tag{3}$$

The RMS amplitude over a window W is:

$$\mathrm{RMS}(t) = \sqrt{\frac{1}{|W|}\sum_{k\in W} x^2(t+k)} \tag{4}$$

Local variance (normalised) is defined as:

$$\mathrm{Var}_{\mathrm{norm}} = \frac{\mathbb{E}[x^2] - (\mathbb{E}[x])^2}{\mathbb{E}[|x|]^2 + \epsilon} \tag{5}$$

where $E[\cdot]$ denotes the sample mean computed within the local analysis window and $\epsilon = 10^{-8}$ is a small positive constant introduced to avoid numerical instability when the denominator approaches zero.

Entropy is computed as:

$$H = -\sum_i p_i \log_2 p_i \tag{6}$$

where p_i denotes the empirical probability of the i-th intensity level within the analysis window, computed as

$$p_i = \frac{n_i}{N}, \tag{7}$$

where n_i is the number of samples assigned to bin i and N is the total number of samples within the window.

Coherence (semblance-based) is defined as:

$$S = \frac{\sum_t (\sum_k x_k(t))^2}{K \sum_t \sum_k x_k^2(t)} \tag{8}$$

All attributes are computed on the seismic grid and co-registered with the amplitude data without additional smoothing. Attribute-specific parameters are kept fixed across all experiments to ensure consistency.

- Instantaneous attributes: Computed using a temporal window of 10 samples.
- Local variance: Computed over a spatial window of 11 traces $\times$ 10 samples.
- Entropy: Computed over a spatial window of 11 traces $\times$ 10 samples using 32 discrete levels.
- Coherence (semblance-based): Computed using a temporal gate of 9 samples, with a stepout of 1 trace in both inline and crossline directions. A cross-shaped analysis window is employed, and a Hamming taper is applied to reduce edge effects.

To complement these formal definitions, Figure 3 presents representative examples of the amplitude and derived attributes after channel-wise normalization. The visual comparison highlights how different attributes emphasise distinct physical characteristics of the seismic signal.

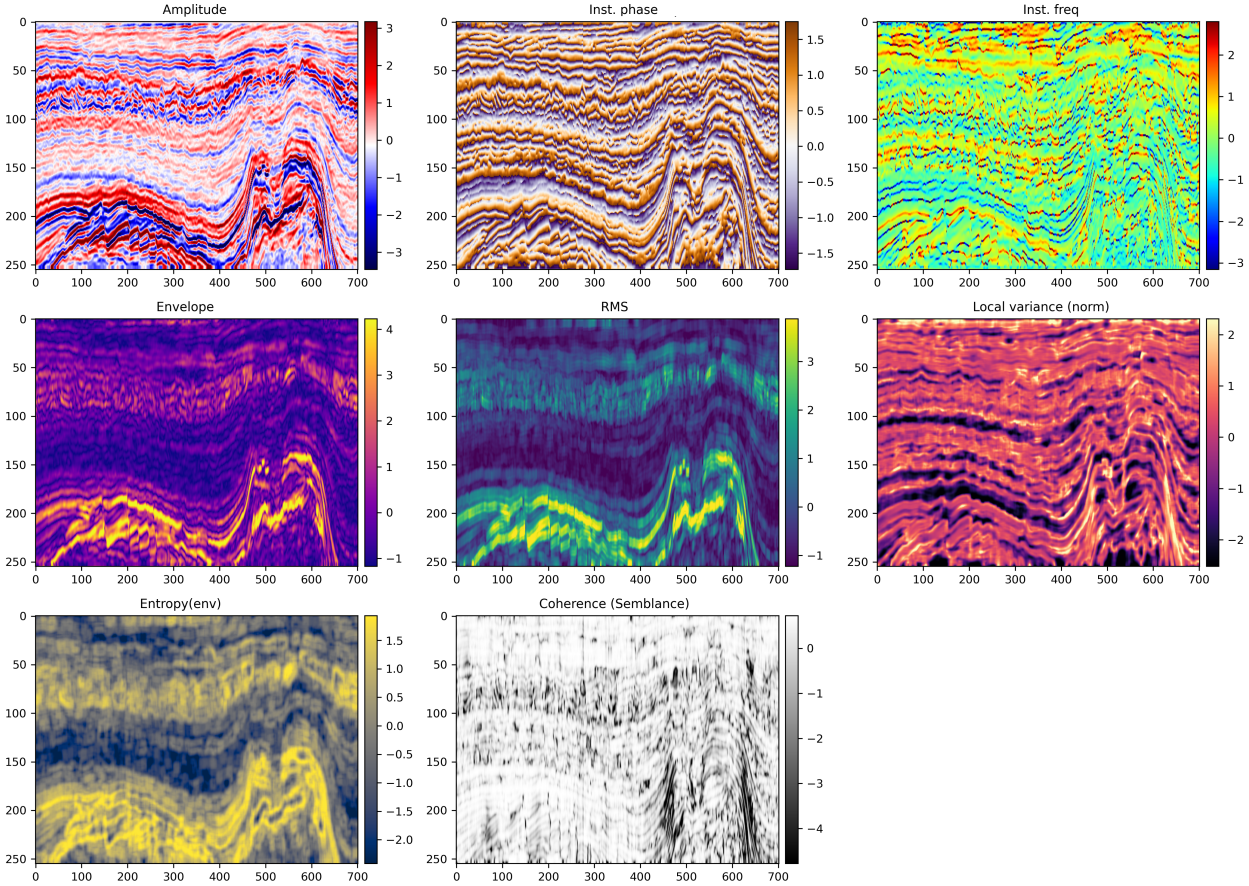

Figure 3. Representative examples of seismic amplitude and derived attributes from the F3 Netherlands dataset after channel-wise normalization. The attributes illustrate different physical properties of the seismic signal: phase (instantaneous phase, instantaneous frequency), energy (envelope, RMS amplitude), texture (local variance, entropy), and structural continuity (coherence). While some attributes exhibit strong visual similarity (e.g., envelope and RMS amplitude), others highlight complementary geological features such as discontinuities and local heterogeneity.

Energy-related attributes, particularly envelope and RMS amplitude, exhibit strong visual similarity, suggesting potential redundancy in the information they encode. Entropy also shares large-scale patterns with energy-based attributes, reflecting its dependence on amplitude distribution. In contrast, structural and texture-related attributes such as coherence and local variance emphasise discontinuities and local heterogeneity, revealing features that are not clearly observable in the amplitude alone. Instantaneous frequency further captures rapid signal variations associated with stratigraphic transitions.

These qualitative observations provide an intuitive interpretation of redundancy and complementarity among attributes, which are quantitatively analysed in subsequent sections through correlation, mutual information, and representation-based similarity measures.

4.2. Input-Space Redundancy Analysis

Statistical redundancy between attributes is evaluated using Pearson correlation, Spearman correlation, and normalized mutual information (NMI). Mutual information is estimated using non-parametric methods [22].

$$r_{ij} = \frac{\mathrm{cov}(X_i, X_j)}{\sigma_{X_i}\sigma_{X_j}} \tag{9}$$

$$\rho_{ij} = \mathrm{corr}(\mathrm{rank}(X_i), \mathrm{rank}(X_j)) \tag{10}$$

$$NMI_{ij} = \frac{2I(X_i; X_j)}{H(X_i) + H(X_j)} \tag{11}$$

where X_i and X_j denote the seismic attributes under comparison, $I(X_i, X_j)$ is the mutual information between attributes X_i and X_j, and $H(X_i)$ and $H(X_j)$ are their corresponding Shannon entropies.

Attributes are considered statistically redundant if:

$$|r_{ij}| > 0.85 \quad \text{or} \quad |\rho_{ij}| > 0.8, \quad \text{and} \quad \mathrm{NMI}_{ij} > 0.6 \tag{12}$$

These thresholds follow standard effect-size interpretations [23]. Hierarchical clustering is applied to identify redundancy groups.

4.3. Controlled Marginal Contribution Experiments

For each attribute A_i, two configurations are defined:

$$C_0 : \text{amplitude-only baseline} \tag{13}$$

$$C_i : C_0 + A_i \tag{14}$$

The marginal contribution is:

$$\Delta\mathrm{mIoU}_i = \mathrm{mIoU}(C_i) - \mathrm{mIoU}(C_0) \tag{15}$$

4.4. Functional Redundancy Metrics

Let $f(X; \theta)$ denote the segmentation model.

Predictive Disagreement (PD)

$$\mathrm{PD}(C_i, C_j) = \frac{1}{N}\sum_{p=1}^{N} \mathbb{I}\left(\hat{y}_{i,p} \neq \hat{y}_{j,p}\right) \tag{16}$$

Error Correlation (EC)

$$\rho_{\mathrm{errors}} = \mathrm{Corr}(E_i, E_j) \tag{17}$$

Centered Kernel Alignment (CKA)

Representational similarity is evaluated using linear Centered Kernel Alignment (CKA) [24], related to Singular Vector Canonical Correlation Analysis (SVCCA) [25].

4.5. Baseline Performance (Amplitude-Only Reference)

The amplitude-only configuration establishes the reference performance for all subsequent experiments. Using the fixed U-Net backbone and the controlled training protocol described in Section 3, the model achieves a test mean Intersection over Union (mIoU) of 0.616, with a pixel accuracy (PA) of 84.52% and a Dice coefficient of 0.832.

Despite relying solely on amplitude, the amplitude-only input captures the dominant stratigraphic structure of the F3 Netherlands dataset. However, this information is not equally discriminative across all facies. While major units such as the Lower North Sea are segmented with high accuracy, more challenging facies such as Scruff and Zechstein exhibit considerably lower IoU values. These results suggest that amplitude alone is insufficient to fully characterise geological variability, particularly in structurally complex or heterogeneous regions.

This uneven behaviour across facies highlights that global metrics alone do not fully describe model performance. In particular, improvements in segmentation are primarily determined by the ability to enhance difficult classes without degrading well-resolved ones. This baseline therefore serves as the reference point for evaluating how different attributes redistribute predictive performance across facies.

4.6. Marginal Contribution Analysis

Table 2 summarises the marginal contribution of each seismic attribute relative to the amplitude-only baseline. Clear performance differences are observed, indicating that not all attributes contribute equally to segmentation accuracy.

Phase-based and structural attributes provide the largest improvements. The instantaneous phase yields the highest gain (ΔmIoU = +0.067), followed by coherence (+0.046) and entropy (+0.028). These attributes introduce complementary information that modifies the decision boundaries of the network, particularly in regions where amplitude alone is insufficient.

In contrast, attributes derived primarily from amplitude energy, such as envelope and RMS amplitude, show limited or negative contribution. RMS amplitude degrades performance (ΔmIoU = −0.006), while envelope also reduces mIoU (ΔmIoU = −0.010), indicating strong redundancy with amplitude and limited additional discriminative power.

Local variance exhibits a different behaviour. While its individual contribution is moderate (ΔmIoU = +0.012), its effect becomes significantly more relevant when combined with phase, as shown in later sections. This suggests that local variance acts as a complementary stabilising attribute, improving segmentation consistency rather than producing strong isolated gains.

Table 2. Marginal contribution of seismic attributes relative to the amplitude-only baseline. Attributes are ranked by ΔmIoU.

Configuration	mIoU	ΔmIoU	Dice	PA
Amplitude + Phase	0.684	+0.067	0.882	88.58%
Amplitude + Coherence	0.663	+0.046	0.847	85.78%
Amplitude + Entropy	0.645	+0.028	0.839	84.94%
Amplitude + Local Variance	0.628	+0.012	0.824	83.88%
Amplitude + Frequency	0.623	+0.006	0.823	83.58%
Amplitude	0.616	—	0.832	84.52%
Amplitude + RMS	0.610	−0.006	0.834	84.30%
Amplitude + Envelope	0.606	−0.010	0.825	82.93%

Overall, these results suggest that attribute usefulness cannot be inferred solely from physical interpretation or statistical distinctiveness, but depends on their ability to alter the model's error patterns in a beneficial manner. Figure 4 illustrates representative segmentation results. While the amplitude-only model captures the overall structure, it struggles in complex regions. The addition of phase improves boundary definition, and the inclusion of local variance further stabilises predictions, reducing inconsistencies in heterogeneous zones.

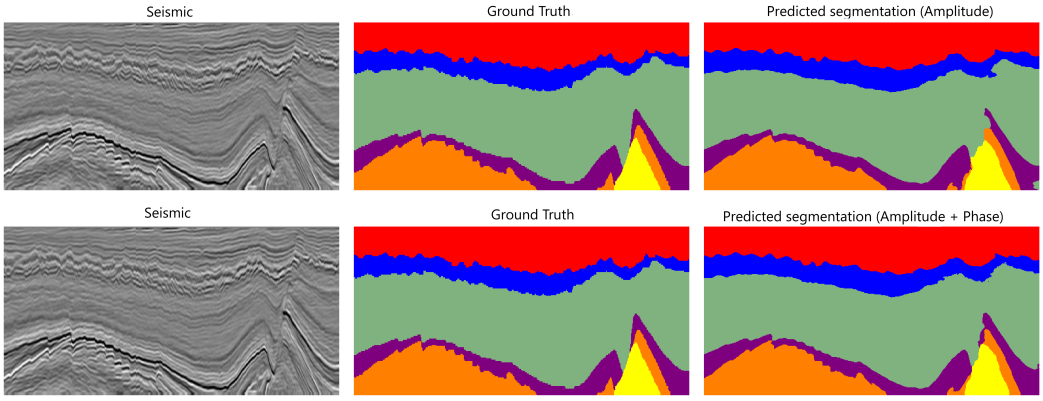

Figure 4. Qualitative comparison of segmentation results. The addition of phase improves boundary delineation and mitigates misclassification in structurally complex regions compared with the amplitude-only baseline segmentation. Each color in the ground-truth and predicted segmentation maps represents one of the six facies classes.

4.7. Per-Class Redistribution Effects

To better understand the impact of seismic attributes beyond global metrics, per-class Intersection over Union (IoU) is analysed for each configuration. The results reveal that attribute contribution is not uniform across facies, but instead redistributes performance among classes.

The amplitude-only baseline already achieves high accuracy in dominant facies such as Lower North Sea, but performs poorly in more challenging units, particularly Scruff and Zechstein. Therefore, improvements in segmentation performance are primarily determined by the ability of attributes to enhance these difficult classes without degrading performance in well-resolved regions.

Instantaneous phase provides the strongest overall improvement. It significantly enhances the segmentation of Upper North Sea and Rijnland/Chalk, and, most importantly, produces a substantial gain in Scruff, one of the most challenging facies. This explains its dominant performance in global metrics and indicates that phase introduces structurally meaningful information that improves decision boundaries.

Coherence exhibits a different but complementary behaviour. While its global improvement is lower than that of phase, it provides the largest gain for the Zechstein facies, which is the most difficult class in the dataset. This suggests that coherence enhances sensitivity to discontinuities and complex structural patterns that are not captured by amplitude or phase alone.

Entropy provides moderate but consistent improvements across multiple facies, including both Scruff and Zechstein, indicating that texture-based information contributes to segmentation without introducing strong trade-offs.

Local variance and instantaneous frequency exhibit non-uniform behaviour. While they improve certain classes, including Zechstein and Rijnland/Chalk, these gains are accompanied by degradation in other facies such as Lower North Sea. This trade-off

explains their limited improvement in global metrics and suggests that these attributes modify the distribution of predictions rather than consistently improving separability.

In contrast, energy-based attributes such as envelope and RMS amplitude degrade performance in the most challenging facies, particularly Zechstein. Although they may improve specific regions, their strong redundancy with amplitude reduces their overall discriminative power.

These results suggest that attribute usefulness cannot be fully characterised by global metrics alone. Instead, attributes redistribute predictive performance across facies, with their effectiveness determined by how they improve difficult classes while preserving performance in dominant ones.

4.8. Input-Space Redundancy Structure

Figure 5 presents the statistical relationships between seismic attributes in the input space, quantified using correlation and normalized mutual information (NMI). Clear redundancy patterns emerge among groups of attributes.

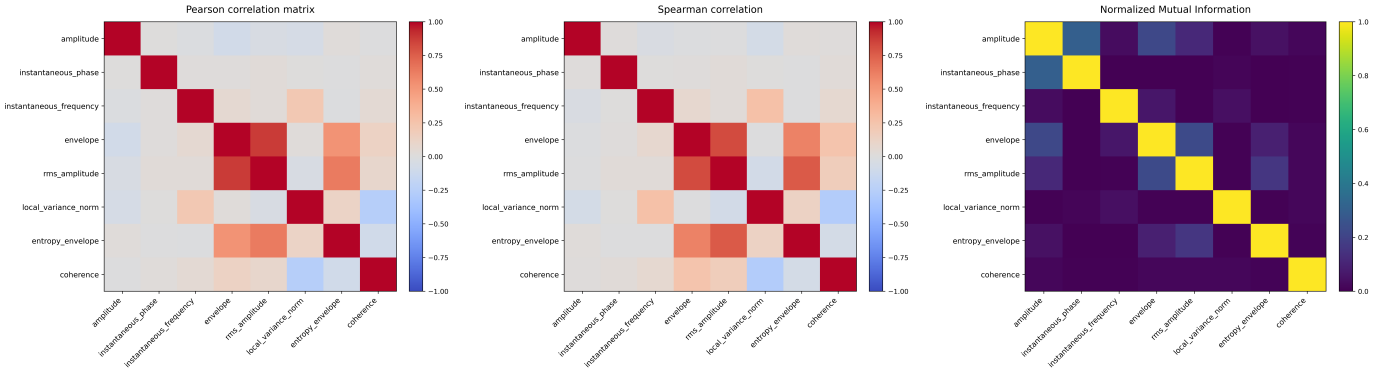

Figure 5. Statistical relationships between seismic attributes. Pearson and Spearman correlations capture linear and monotonic dependencies, while normalized mutual information (NMI) reveals nonlinear relationships.

Energy-related attributes, particularly envelope, RMS amplitude, and entropy, exhibit strong statistical similarity, forming a coherent redundancy group. These attributes capture similar aspects of signal energy and local amplitude distribution, suggesting a high degree of overlapping information.

In contrast, attributes such as local variance and instantaneous frequency exhibit moderate predictive disagreement but relatively high error correlation with the baseline. This indicates that, although their internal representations differ, they tend to produce similar error patterns when used in isolation.

However, this behaviour does not imply that such attributes are functionally irrelevant. In particular, local variance becomes significantly more effective when combined with other attributes, suggesting that its contribution lies in stabilising and refining predictions rather than introducing entirely new decision patterns.

These observations indicate that attributes may influence performance in different ways: some introduce distinct error patterns, while others improve consistency within existing ones.

Hierarchical clustering based on NMI (Figure 6) further confirms this structure, identifying distinct attribute groups corresponding to energy-related attributes and more independent structural and phase-based attributes.

These observations provide a statistical basis for redundancy in the input space. However, as shown in subsequent sections, statistical redundancy alone does not determine

the functional contribution of attributes within deep neural networks, nor does low statistical similarity guarantee improved segmentation performance. In particular, attributes that appear only moderately distinct in the input space may still play a critical role when interacting with other attributes.

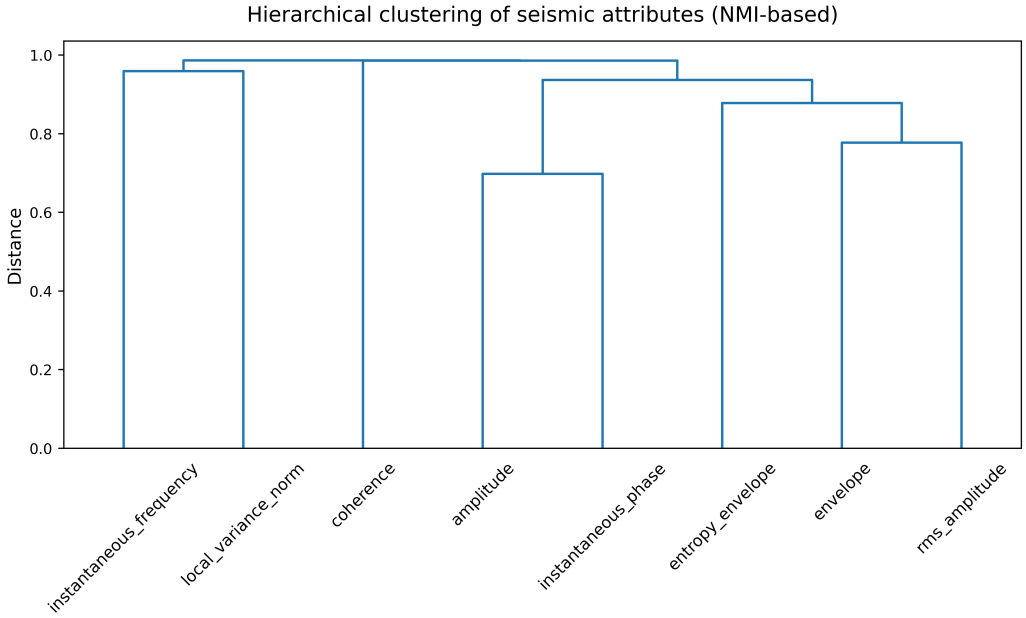

Figure 6. Hierarchical clustering of seismic attributes based on normalized mutual information (NMI). The dendrogram highlights strong redundancy among energy-related attributes (envelope, RMS amplitude, entropy), while attributes such as coherence and phase appear more isolated, indicating lower statistical redundancy.

4.9. Functional Redundancy Analysis

To further investigate redundancy beyond the input space, representational similarity between models is analysed using centered kernel alignment (CKA), together with predictive disagreement (PD) and error correlation (EC). This combined analysis allows us to quantify not only how internal representations differ, but also how these differences translate into segmentation behaviour (Figure 7).

4.9.1. Global Representational Similarity

Figure 7a shows the mean CKA across all model pairs. Overall, representational similarity is high (CKA > 0.96 in most cases), indicating that different attribute configurations converge to similar global representations. This is expected, as all models share the same architecture, task, and amplitude input.

Despite this global similarity, clear relative differences emerge. Energy-related attributes such as envelope and RMS amplitude exhibit extremely high similarity (CKA ≈ 0.998), confirming strong functional redundancy. Entropy also shows high similarity with these attributes, suggesting partial redundancy at the representation level.

In contrast, coherence and instantaneous phase display lower similarity with other configurations, indicating that they induce more distinct internal representations.

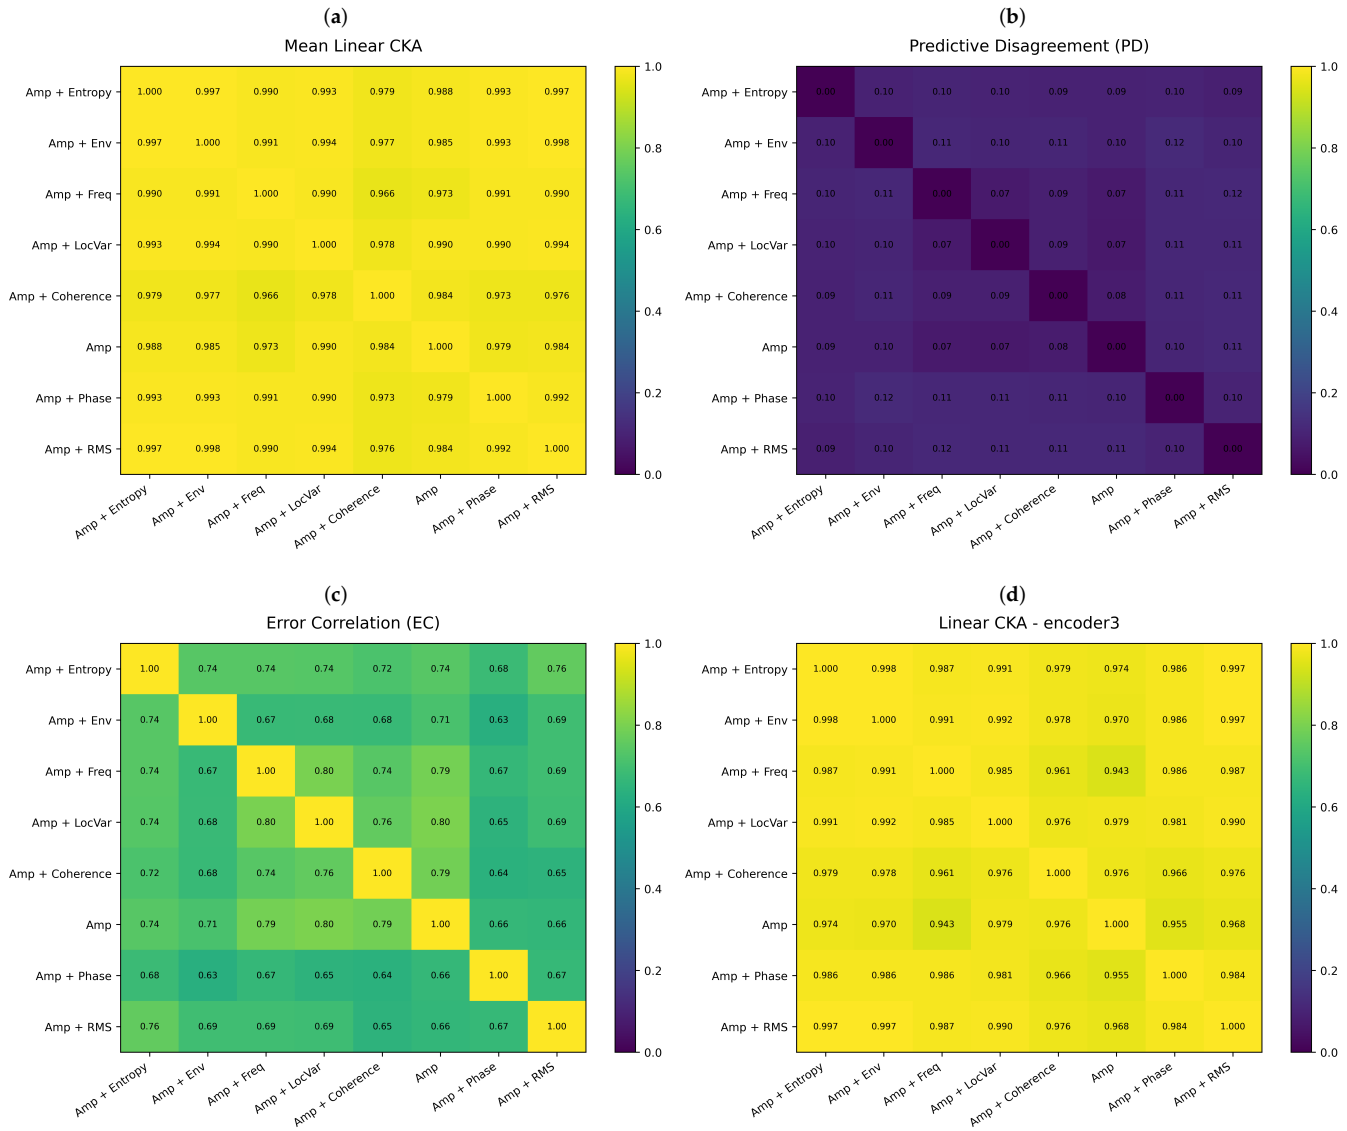

Figure 7. Functional redundancy analysis across attribute configurations. Mean CKA shows high global representational similarity across models (**a**). Predictive disagreement (**b**) and error correlation (**c**) reveal differences in model behaviour and error patterns. Layer-wise CKA in the encoder3 block (**d**) highlights that representational differences are most pronounced in intermediate layers.

4.9.2. Representational Similarity vs. Model Behaviour

Although CKA captures representational similarity, it does not fully explain performance differences. Figure 7b,c show predictive disagreement (PD) and error correlation (EC), respectively, revealing important behavioural differences between models.

Phase exhibits high predictive disagreement with the amplitude-only model and relatively lower error correlation, indicating that it corrects a different subset of errors. This aligns with its strong improvement in segmentation performance.

Coherence also shows complementary behaviour, particularly when compared to phase. While both models share high representational similarity, their error correlation is lower than that of other attribute pairs, suggesting that they contribute differently to segmentation accuracy.

In contrast, attributes such as local variance and instantaneous frequency exhibit moderate predictive disagreement but high error correlation with the baseline. This indicates that, although their internal representations differ, they tend to make similar errors, explaining their limited impact on performance.

Energy-based attributes such as RMS amplitude and envelope show both high representational similarity and high error correlation, confirming their redundancy both statistically and functionally.

4.9.3. Layer-Wise Representation Analysis

Figure 7d presents CKA values at the encoder3 layer. Compared to the global CKA, larger differences are observed at this intermediate level, indicating that attribute-dependent variations are most pronounced during feature extraction.

In deeper layers, particularly at the bottleneck, representations become highly similar across all configurations (CKA > 0.98), suggesting convergence toward a shared semantic encoding of the segmentation task. This indicates that differences introduced by attributes are progressively aligned as the network builds higher-level representations.

Notably, coherence exhibits the largest deviation from other models at intermediate layers, consistent with its role in capturing structural discontinuities. Phase also introduces meaningful variation at this level, contributing to improved segmentation of complex facies.

In contrast, attributes such as local variance exhibit smaller deviations across layers, indicating that their contribution is not driven by large representational shifts, but rather by subtle refinements in feature encoding.

4.10. Redundancy-Aware Subset Validation

Based on the combined analysis of input-space redundancy, per-class performance, and functional similarity, a minimal informative subset of attributes is evaluated. The objective is to determine whether complementary attributes can be combined to improve segmentation performance beyond individual contributions.

4.10.1. Global Performance Analysis

Table 3 summarises the global performance of the evaluated configurations. The amplitude-only baseline achieves an mIoU of 0.616, which is significantly improved by the addition of phase, reaching 0.684 (+0.068). This confirms that phase introduces structural information that substantially enhances segmentation performance.

The combination of amplitude, phase, and local variance achieves the highest overall performance, with an mIoU of 0.697 (+0.081), outperforming both the baseline and the amplitude + phase configuration. In contrast, the addition of coherence (amplitude + phase + coherence) reaches 0.675, which, although higher than the baseline, does not surpass the performance of phase alone.

These results indicate that complementary attributes do not necessarily produce additive improvements. While both local variance and coherence provide additional information, their interaction with phase leads to different outcomes: local variance enhances global performance, whereas coherence introduces interactions that limit further gains.

Table 3. Comparison of segmentation performance across selected attribute configurations.

Configuration	mIoU	ΔmIoU	Dice	PA
Amplitude	0.616	—	0.832	84.52%
Amplitude + Phase	0.684	+0.068	0.882	88.58%
Amplitude + Phase + LocalVar	0.697	+0.081	0.878	88.23%
Amplitude + Phase + Coherence	0.675	+0.058	0.878	88.12%

Although the primary objective of this work is not to benchmark segmentation architectures, it is useful to place the obtained performance into the context of representative CNN-based methods reported on the F3 Netherlands dataset. Table 4 summarizes the best

configuration obtained in this study together with a representative state-of-the-art CNN architecture from the literature. Since different evaluation protocols were adopted, the reported results are provided only for contextual reference and should not be interpreted as a direct architectural comparison.

Table 4. Representative CNN-based semantic segmentation results reported on the F3 Netherlands dataset. Results are provided for performance context only, as different evaluation protocols were adopted across studies.

Method	mIoU (%)	Reference
Improved Deep Dilated CNN (DDCNN) (random split)	72.41	Yang et al. [15]
Proposed framework (official spatial split)	69.73	This work

As shown in Table 4, the proposed redundancy-aware attribute selection framework enables a conventional U-Net to achieve competitive performance without modifying the underlying network architecture. Although Yang et al. [15] achieved a higher mIoU, their experiments were conducted using a random train–test split, whereas the present study adopts the more challenging fixed spatial partition of the F3 Netherlands benchmark. Nevertheless, the relatively small performance gap suggests that careful selection of complementary seismic attributes can substantially improve the performance of a conventional U-Net, bringing it closer to that of considerably more sophisticated CNN architectures.

4.10.2. Per-Class Performance Analysis

To understand the origin of these differences, Table 5 presents the IoU for each facies. The amplitude-only baseline performs well in dominant facies such as Lower North Sea (0.904), but exhibits poor performance in challenging classes such as Scruff (0.412) and Zechstein (0.276).

The addition of phase produces strong improvements in Upper North Sea (0.842) and Scruff (0.602), indicating that phase enhances the delineation of structural boundaries in geologically complex regions. However, this improvement is not uniformly distributed, as performance in Rijnland/Chalk decreases from 0.666 to 0.608.

The inclusion of local variance leads to a different behaviour. Instead of maximising performance in specific facies, it produces more balanced improvements across classes. In particular, it achieves the best performance in Middle North Sea (0.797), Lower North Sea (0.919), Rijnland/Chalk (0.691), and Zechstein (0.438), suggesting that texture-related information is particularly beneficial in geologically heterogeneous facies, while maintaining competitive performance in Upper North Sea and Scruff.

Compared to amplitude + phase, the addition of local variance reduces performance peaks in certain facies (e.g., Upper North Sea and Scruff) but improves weaker classes, resulting in a more balanced distribution of accuracy across the dataset.

Table 5. Per-class IoU comparison for selected configurations.

Configuration	Upper NS	Middle NS	Lower NS	Rijnland/Chalk	Scruff	Zechstein
Amplitude	0.692	0.747	0.904	0.666	0.412	0.276
Amplitude + Phase	0.842	0.767	0.912	0.608	0.602	0.373
Amplitude + Phase + LocalVar	0.776	0.797	0.919	0.691	0.561	0.438

4.10.3. Interpretation

These results suggest two operational forms of functional complementarity under the adopted experimental protocol. Phase exhibits conflict-driven complementarity by

providing targeted improvements that do not always translate into proportional global performance gains. In contrast, local variance exhibits stable complementarity, contributing more balanced improvements across facies, particularly by strengthening weaker classes and reducing performance disparities.

For comparison, coherence improves specific facies such as Zechstein, but its combination with phase does not yield additive global gains, indicating conflict-driven interactions between attributes.

Overall, under the adopted experimental protocol, the combination of phase and local variance appears to represent a form of stable complementarity, where improvements are distributed across classes, leading to the highest global performance. This suggests that the effectiveness of attribute combinations depends not only on the information they provide, but on how this information is integrated within the network.

These findings confirm that attribute selection should be guided by functional complementarity and stability of interaction, rather than by the assumption that additional attributes will yield additive performance gains.

5. Discussion

5.1. Complementarity Does Not Imply Additivity

A central finding of this work is that complementary seismic attributes do not necessarily produce additive improvements when combined. Although attributes such as phase and coherence improve different subsets of errors, their joint use does not yield a proportional increase in global performance.

Importantly, our results suggest that two operational forms of functional complementarity can be distinguished, namely stable complementarity and conflict-driven complementarity, which are formally defined in the following subsection.

This result challenges a common assumption in seismic interpretation workflows, where stacking multiple attributes is expected to systematically improve performance. In deep neural networks, attribute interactions within the feature space and during optimisation can limit the effective integration of complementary information.

Furthermore, statistical similarity in the input space does not necessarily correspond to functional redundancy within the network. Attributes that appear correlated in the input space may still provide distinct and useful contributions during learning.

5.2. Stable vs. Conflict-Driven Complementarity

The results reveal that not all complementary attributes interact in the same manner. Throughout this work, we use the term stable complementarity as an operational interpretive descriptor derived from the joint analysis of the results to describe attribute combinations that consistently produce balanced improvements across classes under the adopted experimental protocol. In contrast, conflict-driven complementarity refers to combinations in which individually informative attributes exhibit competing interactions that limit or destabilise the resulting segmentation performance. Both terms are used as operational interpretive descriptors rather than as independently validated quantitative categories. Under this definition, functional complementarity should not be interpreted as necessarily producing additive performance gains.

Attributes such as phase and coherence exhibit conflict-driven complementarity. Although they improve different subsets of errors, their combination introduces competing interactions that limit global performance gains. This behaviour is reflected in the fact that improvements in specific facies are not consistently translated into overall performance gains.

In contrast, within the experimental setting of this study, local variance appears to exhibit stable complementarity when combined with phase. Rather than producing pronounced gains in specific facies, it contributes to a more balanced distribution of segmentation performance across classes by improving weaker classes while maintaining competitive performance in the remaining ones.

This distinction is consistent with the joint analysis of per-class performance, predictive disagreement, error correlation, and representational similarity, where some attributes modify decision boundaries in complementary ways, whereas others introduce competing effects that limit their joint contribution.

5.3. Error Patterns as the Key Factor

The analysis reveals that attribute usefulness is more closely related to changes in error patterns than to representational differences or statistical independence. Attributes that significantly alter the distribution of errors, such as phase, produce meaningful performance gains, while attributes that modify internal representations without substantially changing error structure have limited impact.

This suggests that evaluating attribute effectiveness requires moving beyond global similarity measures and focusing on how each attribute reshapes the model's decision boundaries.

5.4. Intermediate Representations as the Locus of Variation

Layer-wise analysis shows that differences between attributes are most pronounced in intermediate encoder layers, where local features such as texture, continuity, and structural patterns are extracted. In contrast, deeper layers converge toward similar semantic representations across all configurations.

This indicates that attribute influence is primarily exerted during feature construction rather than at the level of final semantic encoding. As a result, differences in segmentation performance arise from how features are formed, rather than from fundamentally different high-level representations.

5.5. Implications for Attribute Selection

These findings suggest that effective attribute selection should prioritise functional complementarity over diversity in input space. Rather than combining attributes indiscriminately, it is more beneficial to select a minimal set that targets different sources of error.

In this context, attribute selection should consider not only complementarity, but also interaction stability between attributes. While some combinations introduce conflicting effects that limit performance, others provide more consistent and balanced improvements across classes.

For example, phase provides strong structural cues, while local variance contributes to stabilising predictions across heterogeneous regions. Their combination suggests that effective attribute sets can be both complementary and stable, rather than simply diverse.

This suggests that attribute selection should be guided not only by physical interpretation or statistical diversity, but by how attributes interact within the model to modify and stabilise error patterns.

5.6. Limitations and Future Work

The conclusions of this study are derived from a specific dataset, which is relatively limited in size and exhibits class imbalance, as well as from a fixed architecture and training protocol. Repeated training runs with different random initialisations were not performed; therefore, variability across training runs was not assessed. While the proposed framework is generally applicable, the optimal subset of complementary seismic attributes is inher-

ently dataset-dependent and may vary with geological setting, acquisition characteristics, class distribution, and model architecture. Future work will therefore focus on validating the framework across multiple seismic surveys, extending it to 3D seismic volumes, and evaluating its applicability to other seismic segmentation tasks.

6. Conclusions

This work presents a redundancy-aware framework for analysing the contribution of seismic attributes in deep learning-based segmentation. By combining input-space analysis, representational similarity, and error-based evaluation, the proposed approach enables a systematic characterisation of how attributes interact within the model.

The results suggest that seismic attributes can be organised into groups based on statistical similarity in the input space. However, this grouping does not directly translate to functional behaviour within the network. In particular, attributes such as amplitude and instantaneous phase may exhibit statistical similarity while producing distinct error patterns and meaningful performance gains, highlighting a clear discrepancy between input redundancy and functional contribution.

Furthermore, the analysis shows that not all attribute combinations lead to additive improvements. While some attributes provide complementary information, their interaction within the network may introduce conflicts that limit performance gains. At the same time, the results reveal that it is possible to identify minimal informative subsets of attributes that are both complementary and stable.

In particular, combinations that include attributes with different functional roles—such as structural information (phase) and local heterogeneity (local variance)—can improve segmentation performance in a more balanced and stable manner across classes, avoiding the trade-offs observed in other combinations.

Overall, these findings indicate that attribute selection in deep learning-based seismic segmentation should not be guided by input diversity alone, but by functional complementarity and stability of interaction. Rather than identifying a universally optimal attribute combination, the proposed framework provides a principled and transportable methodology for identifying functionally complementary seismic attributes within a given dataset and interpretation task. As such, the framework can be readily extended to other seismic interpretation problems, contributing to more efficient, interpretable, and reliable deep learning workflows.

Author Contributions: Conceptualization, R.C.M.-H. and J.A.M.-H.; methodology, R.C.M.-H., J.A.M.-H. and M.D.l.P.-R.; software, R.C.M.-H.; validation, J.A.M.-H. and M.D.l.P.-R.; formal analysis, R.C.M.-H., J.A.M.-H. and M.D.l.P.-R.; investigation, R.C.M.-H., J.A.M.-H. and M.D.l.P.-R.; resources, J.A.J.-B., D.S., C.d.C.G.-T. and J.G.B.-S.; data curation, R.C.M.-H. and M.D.l.P.-R.; writing—original draft preparation, R.C.M.-H., J.A.M.-H. and M.D.l.P.-R.; writing—review and editing, J.A.M.-H., R.C.M.-H., M.D.l.P.-R., J.A.J.-B., D.S., C.d.C.G.-T. and J.G.B.-S.; visualization, R.C.M.-H. and J.A.M.-H.; supervision, J.A.J.-B., D.S., C.d.C.G.-T. and J.G.B.-S.; project administration, J.A.J.-B., D.S., C.d.C.G.-T. and J.G.B.-S.; funding acquisition, J.A.J.-B., D.S., C.d.C.G.-T. and J.G.B.-S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: An open-source reproducibility repository (MIT License) is available at: https://github.com/RObertoCMH/Redundancy-Aware-Quantitative-Framework (accessed on 21 March 2026). The repository includes the U-Net implementation, training and evaluation scripts,

configuration files (including the random seed), and Python 3.11 code executed in Google Colab using standard scientific libraries to reproduce the figures.

Acknowledgments: The first author acknowledges the National Supercomputing Center and the Instituto Potosino de Investigación Científica y Tecnológica (IPICYT) for hosting the Ph.D. program in Applied Geosciences. The authors also acknowledge the Escuela Superior de Ingeniería Mecánica y Eléctrica (ESIME) and the Escuela Superior de Economía (ESE) of the Instituto Politécnico Nacional. Finally, the authors thank dGB Earth Sciences for making the F3 Netherlands dataset publicly available.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DL	Deep Learning
CNN	Convolutional Neural Network
U-Net	U-shaped Convolutional Network
mIoU	Mean Intersection over Union
IoU	Intersection over Union
PD	Predictive Disagreement
EC	Error Correlation
CKA	Centered Kernel Alignment
SVCCA	Singular Vector Canonical Correlation Analysis
NMI	Normalized Mutual Information
DDCNN	Improved Deep Dilated CNN
RMS	Root Mean Square
PDF	Probability Density Function
Var	Variance
MI	Mutual Information
NNS	Neural Network Similarity
FCN	Fully Convolutional Network
BN	Batch Normalization
ReLU	Rectified Linear Unit
STM	Seismic Time Migration
Post-STM	Post-Stack Time Migration
F3	F3 Netherlands Seismic Dataset

References

1. Bond, C.E.; Lunn, R.J.; Shipton, Z.K.; Lunn, A.D. What makes an expert effective at interpreting seismic images? *Geology* **2012**, *40*, 75–78. [CrossRef]
2. Bond, C.E. Uncertainty in structural interpretation: Lessons to be learnt. *J. Struct. Geol.* **2015**, *74*, 185–200. [CrossRef]
3. Macrae, E.J.; Bond, C.E.; Shipton, Z.K.; Lunn, R.J. Increasing the quality of seismic interpretation. *Interpretation* **2016**, *4*, T395–T402. [CrossRef]
4. Taner, M.T.; Sheriff, R.E. Application of amplitude, frequency, and other attributes to stratigraphic and hydrocarbon determination. In *Seismic Stratigraphy—Applications to Hydrocarbon Exploration*; Payton, C.E., Ed.; AAPG Memoir 26; American Association of Petroleum Geologists: Tulsa, OK, USA, 1977; pp. 301–327. [CrossRef]
5. Chopra, S.; Marfurt, K.J. *Seismic Attributes for Prospect Identification and Reservoir Characterization*; Society of Exploration Geophysicists: Tulsa, OK, USA; European Association of Geoscientists and Engineers: Tulsa, OK, USA, 2007. [CrossRef]
6. Barnes, A.E. *Handbook of Poststack Seismic Attributes*; Society of Exploration Geophysicists: Tulsa, OK, USA, 2016. [CrossRef]
7. Lin, L.; Zhong, Z.; Li, C.; Gorman, A.; Wei, H.; Kuang, Y.; Wen, S.; Cai, Z.; Hao, F. Machine learning for subsurface geological feature identification from seismic data: Methods, datasets, challenges, and opportunities. *Earth Sci. Rev.* **2024**, *257*, 104887. [CrossRef]
8. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional networks for biomedical image segmentation. *arXiv* **2015**, arXiv:1505.04597. [CrossRef]

9. Alaudah, Y.; Michalowicz, P.J.; Alfarraj, M.; AlRegib, G. A multidimensional convolutional neural network for automatic facies identification from seismic data. *arXiv* **2019**, arXiv:1901.07659. [CrossRef]
10. Silva, R.M.; Baroni, L.; Ferreira, R.S.; Civitarese, D.; Szwarcman, D.; Brazil, E.V. Netherlands dataset: A new public dataset for machine learning in seismic interpretation. *arXiv* **2019**, arXiv:1904.00770. [CrossRef]
11. Wrona, T.; Pan, I.; Gawthorpe, R.L.; Fossen, H. Seismic facies analysis using machine learning. *Geophysics* **2018**, *83*, O83–O95. [CrossRef]
12. Qi, J.; Zhang, B.; Lyu, B.; Marfurt, K. Seismic attribute selection for machine-learning-based facies analysis. *Geophysics* **2020**, *85*, O17–O35. [CrossRef]
13. Monteiro, B.A.A.; Canguçu, G.L.; Jorge, L.M.S.; Vareto, R.H.; Oliveira, B.S.; Silva, T.H.; Lima, L.A.; Machado, A.M.C.; Schwartz, W.R.; Vaz-de-Melo, P.O.S. Literature review on deep learning for the segmentation of seismic images. *Earth Sci. Rev.* **2024**, *258*, 104955. [CrossRef]
14. Li, F.; Zhou, H.; Wang, Z.; Wu, X. ADDCNN: An attention-based deep dilated convolutional neural network for seismic facies analysis with interpretable spatial–spectral maps. *IEEE Trans. Geosci. Remote Sens.* **2021**, *59*, 1733–1744. [CrossRef]
15. Yang, N.X.; Li, G.F.; Li, T.H.; Zhao, D.F.; Gu, W.W. An improved deep dilated convolutional neural network for seismic facies interpretation. *Pet. Sci.* **2024**, *21*, 1569–1583. [CrossRef]
16. Alsalmi, H.; Elsheikh, A.H. Automated seismic semantic segmentation using attention U-Net. *Geophysics* **2024**, *89*, WA247–WA263. [CrossRef]
17. Babikir, I.; Abdul Latiff, A.H.; Elsaadany, M.; Pratama, H.; Sajid, M.; Mad Sahad, S.; Ishak, M.A.; Laudon, C. Enhancing machine learning-based seismic facies classification through attribute selection: Application to 3D seismic data from the Malay and Sabah Basins, offshore Malaysia. *Geomech. Geophys. Geo-Energy Geo-Resour.* **2024**, *10*, 138. [CrossRef]
18. Lubo-Robles, D.; Devegowda, D.; Jayaram, V.; Bedle, H.; Marfurt, K.J.; Pranter, M.J. Quantifying the sensitivity of seismic facies classification to seismic attribute selection: An explainable machine-learning study. *Interpretation* **2022**, *10*, SE41–SE69. [CrossRef]
19. Kim, Y.; Hardisty, R.; Marfurt, K.J. Attribute selection in seismic facies classification: Application to a Gulf of Mexico 3D seismic survey and the Barnett Shale. *Interpretation* **2019**, *7*, SE281–SE295. [CrossRef]
20. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2015**, arXiv:1412.6980. [CrossRef]
21. Boashash, B. Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals. *Proc. IEEE* **1992**, *80*, 520–538. [CrossRef]
22. Kraskov, A.; Stögbauer, H.; Grassberger, P. Estimating mutual information. *Phys. Rev.* **2004**, *69*, 066138. [CrossRef] [PubMed]
23. Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed.; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1988.
24. Kornblith, S.; Norouzi, M.; Lee, H.; Hinton, G. Similarity of neural network representations revisited. In *Proceedings of the International Conference on Machine Learning (ICML)*; PMLR: Cambridge, MA, USA, 2019; pp. 3519–3529. Available online: https://proceedings.mlr.press/v97/kornblith19a.html (accessed on 21 March 2026).
25. Raghu, M.; Gilmer, J.; Yosinski, J.; Sohl-Dickstein, J. SVCCA: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. *Adv. Neural Inf. Process. Syst. (NeurIPS)* **2017**, *30*, 6076–6085.