

Learning Stratigraphically Consistent Relative Geologic Time from 3D Seismic Data via Sinusoidal Mapping

Yimin Dou^{1,2}, Xinming Wu^{1,2,*}, Hui Gao^{1,2}, and Zhengfa Bi³

¹State Key Laboratory of Precision Geodesy, School of Earth and Space Sciences, University of Science and Technology of China, Hefei, Anhui, China

²Mengcheng National Geophysical Observatory, University of Science and Technology of China, Mengcheng, 233500, China

³Earth & Environmental Sciences, Energy Geosciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

Abstract. Relative Geologic Time (RGT) estimation from seismic data underpins subsurface structural modeling, depositional analysis, and reservoir characterization, providing the basis for horizon correlation and depositional system reconstruction. Accurate RGT estimation remains challenging because RGT is a topologically constrained continuous field in which local errors readily propagate globally through topological coupling, distorting the overall result. Conventional methods depend heavily on prior information, attribute extraction, and manual interaction, resulting in cumbersome workflows with limited automation. Existing deep-learning approaches predominantly adopt a regression formulation optimized by pixel-wise MSE/-MAE losses, which struggles to recover thin horizons and fails to capture the stratigraphic semantics embedded in the RGT field, yielding limited generalization, unstable stratigraphic ordering, and poor adaptability to diverse structural and depositional settings. We propose RGT-Est, a deep-learning framework that transfers the optimization target from the topologically constrained continuous field into a differentiable sinusoidal space. This representation explicitly encodes the periodic stratigraphic semantics of RGT and alleviates the over-smoothing of fine horizons inherent in direct regression. Pointwise, perceptual, and adversarial losses are jointly imposed in this space to enforce local fidelity, inter-layer consistency, and global structural plausibility, providing both fine-horizon discrimination and global stratigraphic awareness. An optional horizon-guidance module accepts sparse 2D or 3D horizons as priors to satisfy varying precision demands. Trained on synthetic data and evaluated on field seismic surveys featuring dense faulting, large unconformities, steeply dipping strata, folded deformations, and clinoforms, RGT-Est achieves state-of-the-art performance among AI-based methods, and attains substantially higher horizon-correlation accuracy and topological consistency when sparse priors are incorporated. The code and data are available at <https://github.com/douyimin/RGT-Est> and <https://doi.org/10.5281/zenodo.20118902>.

Copyright statement. © Author(s) 2026. This work is distributed under the Creative Commons Attribution 4.0 License.

1 Introduction

Estimating the Relative Geologic Time (RGT) from seismic data is a fundamental task in seismic interpretation, carrying significant scientific and engineering value for subsurface structural modeling, depositional evolution analysis, reservoir characterization, and horizon correlation (Zeng et al., 1998; Stark, 2003, 2004). In an RGT volume, each sample is assigned a relative geologic time value. Under normal stratigraphic conditions without overturned strata or reverse-fault-related structural repetition, RGT generally increases with depth; isosurfaces extracted from the RGT volume correspond to seismic horizons, while faults and unconformities are implicitly represented as discontinuities within the field (Zeng et al., 1998; Stark, 2003). An accurate RGT volume can therefore encode, in a single representation, the horizons, faults, and unconformities present throughout the seismic volume, and has been successfully applied to downstream tasks such as sedimentologic interpretation, geobody detection, implicit structural modeling, and missing well-log prediction (Zeng et al., 2012; Grose et al., 2017; Bader et al., 2018; Karimi and Fomel, 2015; Geng et al., 2020; Yang et al., 2023).

1.1 Conventional RGT Estimation Methods

Conventional approaches to RGT estimation fall broadly into three categories. The first relies on manual horizon picking followed by interpolation: interpreters track as many horizons as possible by hand and interpolate them to construct an RGT volume (Zeng et al., 1998; de Bruin et al., 2007). Although such methods honor the picked horizons exactly, the accuracy of the interpolated regions between horizons is difficult to guarantee, and the manual cost is prohibitive. The second category estimates RGT through seismic instantaneous phase unwrapping (Stark, 2003; Wu and Zhong, 2012), exploiting the temporal information embedded in the instantaneous phase, but is prone to phase jumps and global distortion when the subsurface structure is complex or the signal-to-noise ratio is low. The third category is based on local dip or slope estimation (Fomel, 2002, 2010; Lomask et al., 2006; Wu and Janson, 2017; Wu and Fomel, 2018), in which local reflector orientations are obtained via plane-wave destruction or structure tensors, and a global RGT field is subsequently reconstructed by least-squares fitting. While effective in regions of continuous reflections, these methods rely exclusively on local features and often fail to preserve horizon consistency across faults and unconformities (Figueiredo et al., 2007; Lou et al., 2021). To alleviate this issue, several studies propose first “unfaulting” the seismic image before estimating the RGT (Luo and Hale, 2013; Wu et al., 2016), or inserting manually picked control points on both sides of faults as constraints (Wu and Hale, 2015). However, all such strategies demand substantial prior information and manual interaction, resulting in cumbersome workflows with limited automation.

1.2 Deep-Learning-Based RGT Estimation

In recent years, deep learning—in particular convolutional neural networks (CNNs)—has made substantial progress in seismic data processing (Yu et al., 2018), seismic interpretation (Wu et al., 2019; Gao et al., 2021; Shi et al., 2019), and seismic inversion (Yang and Ma, 2019; Geng et al., 2021), opening new avenues for RGT estimation. Geng et al. (2020) were the first to cast RGT estimation as a regression problem, proposing an end-to-end 2D CNN (a ResNet-50 encoder coupled with a U-Net-style decoder) that predicts RGT directly from a seismic section, optimized with the mean squared error (MSE) loss. Trained

solely on synthetic data, their model generalizes well to field data containing complex faulting and folding, substantially outperforming predictive painting (Wu and Fomel, 2018) and other local-filtering baselines. Building on this direction, Bi et al. (2021) further coupled RGT estimation with fault detection using two separate 3D CNNs to address the two tasks in tandem, while Yang et al. (2023) proposed a Transformer-based Multi-Task Learning (MTL) framework that unifies RGT estimation and fault detection in a single network: the two branches share feature maps and parameters, mutually constrain each other, and additionally accept interpreter-provided horizons as priors to improve robustness under complex fault systems, unconformities, and low-SNR data.

A separate but complementary line of work targets sparse-horizon picking and tracking, in which only a small number of stratigraphic surfaces are extracted from seismic data, typically driven by a few interpreter-provided seed points or sparse manual picks. To reduce the reliance on dense annotations, Wu et al. (2022) proposed a weak-supervised framework that learns variable seismic waveform representations from sparse picks, and Wang et al. (2022) introduced a semi-supervised scheme with virtual adversarial training that propagates a few seed points to dense horizons and iteratively suggests regions lacking control points. To better handle structural discontinuities, Li and Huang (2022) designed a network specifically tailored for tracking horizons across faults, while Liu et al. (2024) fused multiple seismic attributes through an ensemble dense inception transformer to improve robustness in low-SNR or weak-reflection zones. More recently, Zhao and Zhao (2025) adopted a TransUnet backbone that combines convolutional locality with Transformer-based long-range context for higher-fidelity horizon extraction, and Liao et al. (2024) explicitly encoded interpretation uncertainty and a vertical ordering constraint into the training objective, noting that conventional point-by-point losses are ill-suited to horizon tracking and that domain shift further limits generalization to field data. These sparse-horizon methods can yield sharp, interpreter-controlled surfaces in regions where seed points are available, but they operate on one or a few horizons at a time and do not directly produce a globally consistent stratigraphic ordering across the entire volume. RGT estimation and sparse-horizon picking are therefore complementary: the former aims at a dense, volumetric stratigraphic representation, while the latter offers high-precision local control that can naturally serve as priors for the former—an interaction we exploit through the optional horizon-guidance module of RGT-Est.

Despite these advances, existing deep-learning approaches still face several fundamental bottlenecks, which we summarize below.

1.3 Fundamental Limitations of Existing AI Methods

Mismatch between the optimization target and the physical semantics of the RGT field. Existing AI methods predominantly formulate RGT estimation as continuous-field regression with MSE or MAE as the training objective (Geng et al., 2020; Bi et al., 2021). Such pixel-wise losses implicitly assume that prediction errors at neighboring pixels are independent and identically distributed, effectively degrading the RGT field into a set of unrelated scalar predictions. This formulation overlooks the periodic and topological structure that RGT encodes as a “stratigraphic code”—each isoline corresponds to a horizon and the layers follow a strict geologic time order. Under such losses, networks tend to converge to overly smoothed solutions that minimize the aggregate numerical error but fail to capture subtle inter-layer variations, resulting in blurred layer boundaries

and suppressed thin-horizon signals, and thus a poor ability to resolve fine stratigraphic details in complex settings (Yang et al., 2023). In short, MSE/MAE optimization pursues “numerical closeness”, whereas RGT estimation fundamentally requires “stratigraphic correctness”—this mismatch is one of the key reasons why current AI methods struggle to push the accuracy ceiling further.

Inability to guarantee topological consistency. RGT is intrinsically a topologically constrained continuous field: it monotonically increases along depth, its isosurfaces never intersect, and it exhibits geologically specific patterns of discontinuity at faults and unconformities. This global topology implies that small local errors are never isolated: they propagate outward through monotonicity and stratigraphic coupling, eventually breaking the global ordering and producing horizon crossings, jumps, dislocations, or even inversions that are geologically implausible. Pure regression frameworks, however, lack any explicit modeling of such global topology—they can neither penalize topological violations at the loss level nor enforce stratigraphic constraints at the architectural level, but rely entirely on model capacity and data scale to “implicitly learn” the ordering rules. As a result, even small pointwise errors can translate into large-scale departures from the true stratigraphic order, and this implicit learning mechanism is particularly unstable in densely faulted zones, unconformities, and steeply dipping regions (Geng et al., 2020).

Limited generalization and structural adaptability. The two issues above jointly render current AI methods practically unusable for direct RGT estimation from seismic data in production. On the one hand, because the network merely “memorizes” stratigraphic patterns implicitly rather than understanding stratigraphic semantics at the representation level, its generalization depends heavily on the structural coverage of the training data: although performance on synthetic data within the training distribution is often satisfactory, it degrades sharply on unseen structural styles (Geng et al., 2020). On the other hand, the pixel-wise i.i.d. optimization assumption makes the network highly sensitive to input perturbations—noise, amplitude variations, frequency differences—and thus hampers transfer across surveys acquired with different parameters. Consequently, existing AI methods for RGT are mostly demonstrated on carefully curated synthetic datasets or on small, simply structured field slices (Geng et al., 2020; Yang et al., 2023; Bi et al., 2021), while convincing large-scale validation on full 3D field volumes, across different surveys, and in structurally complex regions remains unreported.

1.4 Our Work

To address the above issues, we propose RGT-Est, a deep learning framework that learns relative geologic time from 3D seismic data through a sinusoidal mapping. In contrast to existing regression-based methods, we no longer treat the RGT field as a scalar field to be pixel-wise fitted; instead, the scalar field output by the network is mapped into a differentiable sinusoidal space for modeling and optimization. The sinusoidal space inherently encodes the periodic stratigraphic semantics of RGT: we adopt three sinusoidal channels with linearly decreasing frequencies to build a multi-scale phase encoding of the RGT—the high-frequency channel captures fine-scale horizons such as thin layers, the low-frequency channel encodes large-scale stratigraphic framework, and the combination of the three channels yields a unique and distinguishable representation of any RGT value, thereby simultaneously accommodating local detail and global stratigraphic structure within a unified differentiable framework. In this sinusoidal space, subtle stratigraphic differences are amplified into clearly distinguishable phase variations,

allowing loss functions to operate in a more discriminative feature space. Based on this representation, we jointly impose adversarial, perceptual, and MAE losses in the sinusoidal space: the adversarial loss drives the prediction to approach the real stratigraphic distribution globally, the perceptual loss enforces structural consistency across multi-scale features, and the MAE loss stabilizes pointwise numerical fitting. Together, they balance global topology, structural semantics, and local accuracy, so that the approach captures thin-layer details while retaining a robust global stratigraphic view.

The main contributions of this work are summarized as follows.

1. **A sinusoidal-space modeling paradigm for RGT estimation.** We reformulate RGT estimation from continuous scalar-field regression into a multi-scale phase optimization problem in a sinusoidal space. Three sinusoidal channels with linearly decreasing frequencies explicitly encode the periodic stratigraphic semantics of RGT and yield a unique representation of any RGT value, fundamentally alleviating the over-smoothing of thin layers caused by MSE/MAE losses.
2. **A multi-loss collaborative mechanism for global topological constraints.** We jointly impose adversarial, perceptual, and MAE losses in the sinusoidal space, constraining the network from three complementary perspectives—distributional consistency, structural fidelity, and pointwise accuracy—and equipping it with both fine-horizon discrimination and robust global stratigraphic awareness.
3. **Optional sparse horizon guidance.** We design an optional Horizon Guidance module that accepts sparse 2D or 3D horizons as priors. RGT-Est operates fully automatically without any prior; once horizons are provided, it delivers substantially higher precision and naturally preserves lateral consistency in slice-by-slice 3D prediction.
4. **Systematic multi-scenario generalization evaluation.** We evaluate RGT-Est on multiple structurally complex field seismic datasets covering unconformities, densely faulted systems, steeply dipping structures, and strong structural superposition. RGT-Est substantially outperforms publicly available AI-based RGT estimation method—producing high-quality RGT fields end-to-end in direct inference mode, and reaching very high precision once sparse horizons are provided—validating its strong generalization across surveys, structural styles, and data qualities.

2 Method

2.1 Overview

The overall framework of RGT-Est is illustrated in Fig. 1. The input consists of a 3D seismic volume $\mathbf{S} \in \mathbb{R}^{1 \times D \times H \times W}$ together with an optional sparse horizon guidance $\mathbf{H} \in \mathbb{R}^{1 \times D \times H \times W}$ (2D, 3D, or empty), accompanied by an indicator mask $\mathbf{M} = \mathbb{I}(\mathbf{H} \neq 0)$ that marks the positions of user-provided horizons. When no prior is available, $\mathbf{H}$ and $\mathbf{M}$ are simply set to zero, so that the method naturally supports both fully automatic inference and sparse-horizon-guided modes within a single architecture. A 3D HRNet serves as the backbone G_θ and predicts, in an end-to-end manner, a continuous RGT scalar field $\hat{R} \in [-1, 1]^{1 \times D \times H \times W}$ through

$$\hat{R} = \tanh(G_\theta(\mathbf{S}, \mathbf{H}, \mathbf{M})). \quad (1)$$




Figure 1. Overview of the proposed RGT-Est framework. A 3D HRNet backbone takes a 3D seismic volume with optional 2D / 3D horizon guidance as input and predicts a continuous RGT scalar field $\hat{R}$. The Sinusoidal Mapping module maps $\hat{R}$ into a three-channel phase space via $\sin(f_i \hat{R})$ with $f_1=2.0$, $f_2=1.0$, $f_3=0.5$, yielding the Sinusoidal RGT. Training jointly minimizes the MAE, perceptual, and adversarial losses over the Predicted RGT and the Sinusoidal RGT representations.

The prediction $\hat{R}$ is then fed into the *Sinusoidal Mapping* module $\mathcal{T}$, which maps the scalar field into a three-channel phase space via frequency-decreasing sinusoidal encodings; the resulting Sinusoidal RGT amplifies subtle stratigraphic differences and explicitly expresses the periodic stratigraphic semantics embedded in $\hat{R}$. A 3D PatchGAN discriminator D_ϕ is introduced to distinguish the predicted phase representation from the ground-truth one, conditioned on the input seismic volume. During training, three complementary losses—a hybrid regression loss, a multi-directional perceptual loss, and a relativistic adversarial loss—are jointly minimized to reconcile local accuracy, structural semantics, and global topology.

2.2 Sinusoidal Mapping

A key innovation of RGT-Est is that losses are *not* evaluated solely on the scalar RGT field $\hat{R}$; instead, $\hat{R}$ is first lifted into a higher-dimensional phase space through a differentiable Sinusoidal Mapping module $\mathcal{T}$. Given $\hat{R} \in [-1, 1]$, we first rescale it to the discretized range $[1, N]$ with $N = 256$:

$$R = \frac{\hat{R} + 1}{2} (N - 1) + 1, \quad (2)$$

and then construct a three-channel phase representation using sinusoidal basis functions with linearly decreasing frequencies:

$$\mathcal{T}(R) = [\sin(f_1 R), \sin(f_2 R), \sin(f_3 R)]^\top, \quad (f_1, f_2, f_3) = (2.0, 1.0, 0.5). \quad (3)$$

The geometric frequency decay ($f_1 : f_2 : f_3 = 4 : 2 : 1$) is chosen so that the three channels jointly span a wide range of stratigraphic scales: the high-frequency channel $\sin(f_1 R)$ captures thin layers and fine-scale horizons, the low-frequency channel $\sin(f_3 R)$ encodes the large-scale stratigraphic framework, and the mid-frequency channel interpolates between the two. In this

multi-frequency space every RGT value is mapped to a unique, distinguishable three-dimensional phase vector, preventing the wrap-around ambiguity that would arise from a single-frequency encoding. Moreover, minute variations in R are amplified into clearly separable phase differences in $\mathcal{T}(R)$, allowing downstream losses to operate in a more discriminative feature space. Throughout the remainder of the paper, we denote the predicted and ground-truth phase representations by

$$\hat{\mathbf{P}} = \mathcal{T}(\hat{R}), \quad \mathbf{P} = \mathcal{T}(R). \quad (4)$$

To empirically verify the mechanism behind the proposed Sinusoidal Mapping, we perform a gradient back-propagation analysis that compares the training signal received by the prediction under the raw scalar representation versus the sinusoidal phase representation (Fig. 2). We freeze a pretrained 3D HRNet, feed a real seismic volume, and take its forward output $\hat{R}$ as the prediction; the ground-truth RGT R is retained as the regression target. Two losses—the L_1 loss and the multi-directional 3D LPIPS loss—are computed in two equivalent domains: (i) directly on the scalar RGT field, and (ii) on the three-channel phase field $\mathcal{T}(\cdot) = [\sin(f_1 R), \sin(f_2 R), \sin(f_3 R)]$ with $(f_1, f_2, f_3) = (2.0, 1.0, 0.5)$ as defined in Sect. 2.2. For every configuration the backward pass is traced to the prediction, and the gradient tensor at $\hat{R}$ is recorded and compared.

Figure 2a contrasts the two domains on a central inline slice. The raw-space L_1 gradient collapses to a nearly piecewise-constant sign map ($|g|_{\text{mean}} = 1.73 \times 10^{-8}$), because the pixelwise $\mathcal{L}_{L_1} = |\hat{R} - R|$ produces a gradient of magnitude $1/N$ that is completely blind to whether the error occurs at a thin horizon or a homogeneous interior. More revealingly, in panel (a) this raw-space L_1 gradient exhibits a sharp positive–negative interface that partitions the slice into two nearly uniform regions, indicating that the regression signal is merely pulling $\hat{R}$ toward or away from the bulk depth trend rather than encoding any stratigraphic structure; the raw-space LPIPS gradient suffers from the same limitation to a milder extent, responding only to a few high-amplitude reflectors and leaving the vast majority of layer interiors essentially unsupervised. In sharp contrast, the sinusoidal-space gradients of both L_1 and LPIPS exhibit strong, spatially coherent responses that trace the underlying stratigraphic layering—exactly the fine-scale horizons that conventional regression losses tend to smooth out. The gradient magnitude increases by orders of magnitude, confirming that the same prediction error produces a substantially stronger and more discriminative training signal once it is measured in the sinusoidal phase space.

Figure 2b,c plot the per-voxel distributions of $|\nabla|$ on a logarithmic scale. The raw-space distributions are narrowly concentrated near 10^{-8} – 10^{-7} with a heavy truncation at the lower tail (the gradient magnitude is bounded by $1/N$ for L_1). The sinusoidal-space distributions are shifted by roughly one-to-two decades to the right and exhibit a much longer tail, with medians moving from 1.7×10^{-8} to 1.0×10^{-6} for L_1 and from 1.3×10^{-7} to 2.6×10^{-6} for LPIPS. A richer tail indicates that more voxels receive informative, non-saturating gradient updates, which translates into finer-grained parameter updates during training.

Figure 2d,e show the root-mean-square of $|\nabla|$ computed per depth slice. For raw-space L_1 , the profile is an almost flat vertical line because the loss degenerates to sign-based updates and has no depth-dependent structure. For the sinusoidal-space counterpart, the profile exhibits clear stratigraphic modulation—larger gradients at layer boundaries and smaller gradients in between—which is precisely the behavior required to correct thin-horizon errors without over-smoothing the interior. The LPIPS profiles show the same trend but at a higher baseline.




Figure 2. Gradient back-propagation comparison between the raw RGT space and the proposed sinusoidal space. A pretrained 3D HRNet is kept frozen, and the L_1 and 3D LPIPS losses are computed either directly on $\hat{R}$ or on its sinusoidal encoding $\mathcal{T}(\hat{R})$; the gradient at $\hat{R}$ is recorded for comparison. (a) Inline slice of the seismic input, the GT RGT, the prediction, and the resulting L_1 /LPIPS gradient maps in both domains. (b,c) Log-scale per-voxel distributions of $|\nabla|$. (d,e) Depth-wise RMS of $|\nabla|$. Right: summary statistics and the overall amplification ratio $|\nabla|_{\text{sin}}/|\nabla|_{\text{raw}}$ ($62.1\times$ for L_1 , $15.8\times$ for LPIPS).

Across the volume, computing the loss in the sinusoidal space amplifies the mean gradient magnitude received by the prediction by $62.1\times$ for L_1 and $15.8\times$ for LPIPS (right panel of Fig. 2). This is consistent with the first-order estimate

$$\frac{\partial \mathcal{T}_c(R)}{\partial R} = f_c \cdot \frac{N-1}{2} \cdot \cos(f_c s), \quad (5)$$

where $s \in [1, N]$ is the rescaled RGT value and $N=256$: the product $f_c \cdot (N - 1)/2$ already yields an amplification factor on the order of 10^2 per channel, which is then modulated by the local phase term $\cos(f_c s)$ and redistributes the energy toward layer boundaries. The sinusoidal representation therefore acts as a *differentiable, stratigraphy-aware amplifier* of the training signal, explaining why the multi-objective losses defined in Sect. 2.5 are particularly effective at recovering thin horizons and preserving stratigraphic ordering.

2.3 3D HRNet Backbone

We adopt a three-dimensional High-Resolution Network (HRNet) as the backbone G_θ because its parallel multi-resolution branches are well suited for RGT estimation: the high-resolution branch preserves fine-scale horizon details, while the low-resolution branches enlarge the receptive field so that global stratigraphic structure can be captured. Given the concatenated input $[S, H, M]$ of three channels, a two-layer strided convolution stem first downsamples the volume by a factor of four. Stage 1 is a sequence of four Bottleneck residual blocks operating at $1/4$ resolution with 256 channels; Transition 1 then branches out an additional $1/8$ -resolution stream with $2c$ channels ($c = 48$ in all experiments). Stage 2 stacks two Stage-Modules, each containing four BasicBlock residual units per branch and a cross-resolution fusion that exchanges information between the two parallel streams. Transition 2 further introduces a $1/16$ -resolution branch with $4c$ channels, and Stage 3 contains six StageModules that repeatedly fuse the three parallel streams through learnable upsampling, downsampling, and identity connections. After Stage 3 the highest two resolutions ($1/4$ and $1/8$) are retained, concatenated after bilinear upsampling, and fed into a regression decoder consisting of $3 \times 3 \times 3$ convolutions interleaved with two PixelShuffle3D layers that progressively recover the original spatial resolution. A final tanh activation constrains the predicted RGT to $[-1, 1]$, matching the normalization range of the training labels and stabilizing the subsequent sinusoidal transform.

2.4 Discriminator Architecture

To impose a distribution-level topological constraint on the predicted RGT, we employ a 3D PatchGAN discriminator (Isola et al., 2017) D_ϕ conditioned on the input seismic volume. D_ϕ takes as input a four-channel volume formed by concatenating the phase representation (three channels) with the seismic volume (one channel), and outputs a dense patch-wise logit map. Concretely, D_ϕ consists of three strided $4 \times 4 \times 4$ convolutional blocks with base width $b = 100$ and a stride of 2, each (except the first) followed by 3D instance normalization and a LeakyReLU activation with negative slope 0.2, and is terminated by a $4 \times 4 \times 4$ convolution with unit stride that produces the patch logits. The patch-wise design allows D_ϕ to judge local stratigraphic plausibility across the entire volume, rather than committing to a single global decision.

2.5 Loss Functions

RGT-Est is trained by minimizing a weighted combination of three complementary losses operating on both the scalar field $\hat{R}$ and the phase representation $\hat{P}$.

Hybrid regression loss

A hybrid L_1/L_2 loss ℓ is first applied on the scalar RGT field to enforce pointwise numerical fidelity:

$$\ell(a, b) = \|a - b\|_1 + 0.5 \|a - b\|_2^2. \quad (6)$$

Building on this, a regression loss that couples the scalar and phase domains is defined as

$$\mathcal{L}_{\text{reg}} = 2\ell(\hat{R}, R) + 0.2 \|\hat{\mathbf{P}} - \mathbf{P}\|_1. \quad (7)$$

The scalar term stabilizes large-scale fitting while the phase term amplifies the penalty on fine-scale stratigraphic errors, which are hard to recover from pure scalar regression.

Multi-directional 3D perceptual loss

To enforce multi-scale structural consistency, we extend the 2D LPIPS (Zhang et al., 2018) metric to three dimensions by means of a *multi-directional slicing* strategy. For any two volumes $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{N \times C \times D \times H \times W}$, we uniformly sample $k = 32$ slice indices along each axis and compute the 2D LPIPS similarity on the extracted slices, which are then averaged across all three directions:

$$\mathcal{L}_{\text{per}}^{3D}(\mathbf{X}, \mathbf{Y}) = \frac{1}{3} \sum_{d \in \{D, H, W\}} \frac{1}{k} \sum_{i \in \mathcal{I}_d} \text{LPIPS}_{\text{AlexNet}}(\mathbf{X}_i^{(d)}, \mathbf{Y}_i^{(d)}), \quad (8)$$

where $\mathcal{I}_d$ denotes the sampled indices along axis d and $\mathbf{X}_i^{(d)}$ denotes the i -th 2D slice of $\mathbf{X}$ along axis d . The overall perceptual loss is the sum of the perceptual distances computed on the phase representation and on the scalar field:

$$\mathcal{L}_{\text{lips}} = \mathcal{L}_{\text{per}}^{3D}(\hat{\mathbf{P}}, \mathbf{P}) + \mathcal{L}_{\text{per}}^{3D}(\hat{R}, R). \quad (9)$$

This formulation encourages the network to match the ground truth not only pointwise but also in terms of high-level perceptual structure along all three spatial axes.

Relativistic hinge adversarial loss

To impose a distribution-level topological constraint, we adopt the relativistic hinge adversarial formulation of Jolicoeur-Martineau (2018), which combines the hinge GAN loss (Lim and Ye, 2017) with the relativistic discriminator scheme. Let $\hat{\mathbf{P}}$ and $\mathbf{P}$ denote the predicted and ground-truth phase representations, both conditioned on the seismic input $\mathbf{S}$. Writing $d_r = D_\phi([\mathbf{P}, \mathbf{S}])$ and $d_f = D_\phi([\hat{\mathbf{P}}, \mathbf{S}])$, the discriminator loss is

$$\mathcal{L}_D = \frac{1}{2} \mathbb{E}[\text{ReLU}(1 - (d_r - \mathbb{E}[d_f]))] + \frac{1}{2} \mathbb{E}[\text{ReLU}(1 + (d_f - \mathbb{E}[d_r]))], \quad (10)$$

and the generator adversarial loss is correspondingly

$$\mathcal{L}_{\text{adv}} = \frac{1}{2} \mathbb{E}[\text{ReLU}(1 + (d_r - \mathbb{E}[d_f]))] + \frac{1}{2} \mathbb{E}[\text{ReLU}(1 - (d_f - \mathbb{E}[d_r]))]. \quad (11)$$

Relativistic losses compare each real/fake sample against the expected opposite, which is known to yield more informative gradients and more stable training than the non-saturating or hinge variants. Because $\mathcal{L}_{\text{adv}}$ is computed in the phase space, the discriminator effectively regularizes the *stratigraphic distribution* of the prediction rather than its raw scalar values.

Total generator objective

The generator is trained with the weighted sum

$$\mathcal{L}_G = \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{lpips}} \mathcal{L}_{\text{lpips}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}, \quad (12)$$

with $\lambda_{\text{reg}} = 5$, $\lambda_{\text{lpips}} = 5$, and $\lambda_{\text{adv}} = 0.01$. The discriminator is trained with $\lambda_{\text{adv}} \mathcal{L}_D$ using the same adversarial weight.

2.6 Training Data

To train RGT-Est, we randomly generated 2,000 pairs of synthetic 3D volumes (Wu et al., 2020), each with a size of $512 \times 512 \times 512$ (an example is shown in Fig. 1). Each pair consists of a seismic amplitude volume and its corresponding ground-truth RGT volume. The synthetic models were designed to cover a wide range of structural and stratigraphic scenarios that commonly challenge RGT estimation, including normal and reverse faults, densely faulted systems, unconformities, complex folds, slope structures, and laterally varying depositional architectures. These geological elements introduce reflector discontinuities, abrupt stratigraphic terminations, strong lateral variations in dip, and multi-scale deformation patterns, thereby providing diverse training examples for learning stratigraphically consistent RGT representations. The corresponding seismic volumes were generated from the synthetic RGT models through forward modeling, so that the seismic reflections are consistent with the underlying stratigraphic ordering.

2.7 Optimization and Implementation

The generator G_θ and discriminator D_ϕ are optimized alternately in a two-player minimax game. In each iteration, we first update G_θ by minimizing $\mathcal{L}_G$ with D_ϕ 's parameters frozen, and then update D_ϕ by minimizing $\lambda_{\text{adv}} \mathcal{L}_D$ on detached generator outputs. Both networks are optimized with the Adam optimizer (Kingma and Ba, 2015) using $\beta_1 = 0$, $\beta_2 = 0.9$, and a constant learning rate of 10^{-4} , a configuration empirically shown to stabilize relativistic-hinge GAN training. To enlarge data diversity at negligible cost, at each training step we randomly swap the last two spatial axes of the seismic, target RGT, and horizon-guidance volumes with probability 0.5, which acts as an isotropic in-plane augmentation. All experiments are implemented in PyTorch Lightning with mixed-precision training on NVIDIA A100 80G $\times$ 4 GPUs. Batch size is set to 4, and the training input dimension is $400 \times 512 \times 256$ (timeline, inline, crossline).

3 Results

We conduct three sets of experiments to evaluate RGT-Est. **First**, we isolate the contribution of the sinusoidal mapping by directly comparing optimization in the proposed phase space with optimization in the raw RGT voxel space—the formula-

tion used by DeepRGT (Bi et al., 2021) and representative of existing AI-based RGT estimation methods. To ensure a fair comparison, the backbone and training data are held identical across the two settings. **Second**, we examine how sparse 2D (single-section) and 3D (volumetric) horizon priors progressively improve inference accuracy in structurally complex regions. **Third**, we test the generalization of RGT-Est on multiple highly challenging field surveys featuring densely faulted zones, large unconformities, steeply dipping strata, folded deformations, and clinoforms.

3.1 Comparison with Voxel-Space Optimization

To assess the contribution of the proposed sinusoidal mapping, we compare RGT-Est against a re-implemented version of DeepRGT (Bi et al., 2021), the most representative AI-based RGT estimation method that performs optimization directly in the raw RGT voxel space. To remove confounding factors arising from differences in network capacity or training data, we re-implement DeepRGT by replacing its original backbone with the same 3D HRNet used in this work and retraining it on the same training set as RGT-Est, while keeping the rest of its formulation—namely the voxel-space regression loss, optimization schedule, augmentation strategy, and inference protocol—unchanged. Hereafter we refer to this baseline as DeepRGT[†] to indicate that it shares the backbone and training data with RGT-Est. We evaluate both methods on four field surveys with relatively mild structural deformation: the Costa Rica fore-arc basin and three publicly available Netherlands datasets. These surveys span a range of acquisition geometries and stratigraphic styles while remaining within a regime where pixelwise regression baselines are expected to behave reliably, thereby providing a fair benchmark for examining the additional value brought by sinusoidal-space optimization. For each survey, we present representative inline and time-slice visualizations that highlight the differences in fine-scale layering and global structural coherence between the two methods.

Figure 3 presents the field-data comparison between DeepRGT[†] and the proposed RGT-Est on the Costa Rica and Netherlands datasets. The first row shows the Costa Rica data, whereas the remaining three rows show examples from the Netherlands data. For each example, the first column shows the input seismic volume, the second column shows the result of DeepRGT[†], and the third column shows the result of RGT-Est. Because reliable RGT estimation requires both local horizon fidelity and global stratigraphic ordering, we mainly evaluate the results from four aspects: alignment between RGT contours and seismic reflectors, preservation of thin layers, continuity across laterally varying structures, and topological consistency in complex deformation zones.

In the Costa Rica example shown in the first row, the seismic image contains weak reflectors, discontinuous events, and strong structural deformation. These features make RGT estimation difficult because local errors can easily propagate into large-scale distortions of the stratigraphic field. DeepRGT[†] recovers the general depth-increasing trend, but the estimated contours show clear instability in the highlighted regions. In the blue box, where the seismic reflections are relatively weak and closely spaced, the contours produced by DeepRGT[†] are over-smoothed and fail to sufficiently resolve fine stratigraphic variations. In the red box, where the reflectors are strongly deformed, the contours become locally distorted and are less consistent with the seismic events. The magenta box further shows that DeepRGT[†] has difficulty maintaining stable layer continuity in regions with lateral structural changes. By contrast, RGT-Est produces contours that are more consistent with the reflector geometry. The predicted RGT field better follows the deformed seismic events, preserves more detailed stratigraphic variations, and maintains a more




Figure 3. Qualitative comparison of RGT estimation on field seismic volumes. From left to right, each group presents the input seismic volume, the result of DeepRGT[†] (re-implementation) (Bi et al., 2021), and the result of the proposed RGT-Est. For both methods, the displayed horizons are extracted as iso-surfaces from the estimated RGT fields and overlaid on the seismic volumes for visual comparison. The colored dashed boxes highlight challenging regions, including weak reflectors, slope structures, faulted zones, and laterally varying stratigraphy. RGT-Est better follows seismic reflectors and preserves more coherent stratigraphic ordering.

coherent ordering from shallow to deep intervals. This result indicates that RGT-Est generalizes well to structurally complex field data with weak and discontinuous reflectors.

The second row shows the Netherlands F3 dataset, where a prominent slope structure is present in the shallow interval, as highlighted by the blue box. This case is challenging because the reflectors within the slope are laterally inclined and locally thinned, requiring the estimated RGT contours to follow both the dipping geometry and the internal stratigraphic order.

DeepRGT[†] captures the overall stratigraphic trend, but its contours in the slope zone are relatively over-smoothed and do not fully conform to the inclined seismic reflectors. As a result, the internal layering of the slope is weakened, and the lateral variation of the depositional architecture is not sufficiently preserved. In the deeper red-box region, where the reflectors become weaker and more structurally complicated, DeepRGT[†] also produces locally unstable contours. In comparison, RGT-Est better preserves the internal geometry of the slope. Its contours remain more parallel to the inclined reflectors and show stronger lateral continuity across the shallow depositional structure. In the deeper complex region, RGT-Est also gives a more coherent stratigraphic field, suggesting improved robustness to weak reflections and local structural disturbance.

The third row further evaluates the methods on a Netherlands example with stronger structural variation and multiple fault-related discontinuities. In the input seismic volume, the blue, red, and magenta boxes mark regions with different interpretation difficulties, including weakly expressed strata, faulted or steeply varying reflectors, and laterally changing seismic patterns. DeepRGT[†] gives a reasonable large-scale RGT trend, but its contours become less reliable in these complex regions. In particular, around the central faulted zone marked by the red box, the contours tend to be locally blurred and poorly aligned with the displaced reflectors. In the magenta-box region, the estimated contours are also less consistent with the lateral changes in reflector geometry. These artifacts suggest that direct regression tends to suppress subtle stratigraphic differences and may produce locally inconsistent RGT fields when the seismic structures depart from simple layered patterns. RGT-Est substantially improves the result. The predicted contours better track the seismic reflectors on both sides of the faulted zone and preserve clearer stratigraphic separation across laterally varying structures. The result shows fewer local distortions and a more stable ordering of the RGT field, demonstrating the advantage of the proposed stratigraphy-aware optimization.

The fourth row shows another Netherlands example characterized by broad weak-reflection intervals and local deformation. The blue box covers a large region where the reflectors are relatively subtle and laterally continuous, whereas the red box highlights a structurally complicated zone with rapid changes in reflector geometry. DeepRGT[†] tends to produce overly smooth contours in the weak-reflection interval, which reduces the resolution of fine stratigraphic layers. Near the red-box region, its contours also show local distortion and insufficient conformity with the seismic events. RGT-Est produces a more geologically plausible RGT field. In the blue-box region, the contours are more continuous and better aligned with the subtle reflectors, indicating improved sensitivity to weak but coherent stratigraphic signals. Around the red-box region, RGT-Est better adapts to the local deformation and preserves a more consistent relationship between adjacent layers.

Overall, the field-data comparisons demonstrate that RGT-Est consistently outperforms DeepRGT[†] across different surveys and structural settings. DeepRGT[†] can recover the first-order stratigraphic trend, but it often suffers from over-smoothing, weakened thin-layer discrimination, and local contour distortion in weak-reflection zones, slope structures, faulted regions, and laterally varying depositional architectures. In contrast, RGT-Est produces RGT contours that are more consistent with seismic reflectors and more stable across complex structural transitions. The improvement is particularly evident in the highlighted regions, where the proposed method better preserves fine stratigraphic details while maintaining global topological consistency. These observations support the effectiveness of transferring RGT optimization from the raw scalar space to the sinusoidal phase space, where subtle inter-layer differences can be more explicitly represented and jointly constrained by pointwise, perceptual, and adversarial losses.




Figure 4. Effect of stratigraphic constraints on RGT estimation. The purple lines represent the input 2D horizon constraints, which also serve as reference lines for the horizons. (a) Incorporating 2D horizon constraints into RGT-Est. The purple dashed curves denote ground-truth horizons used for visual comparison, and the yellow boxes highlight regions where the constrained result better honors the target stratigraphic geometry. (b) Incorporating sparse 3D horizon constraints into RGT-Est. The input seismic volume and 3D horizon constraints are shown on the left, while the direct inference result and the constraint-guided result are shown in the middle and on the right, respectively. Both 2D and 3D constraints significantly improve horizon alignment and spatial consistency of the estimated RGT field.

3.2 Stratigraphic-Constraint-Guided RGT Estimation

Figure 4 evaluates the effect of incorporating sparse stratigraphic constraints into RGT-Est. The experiment demonstrates that the proposed framework supports both direct inference and constraint-guided inference, allowing user-provided 2D or 3D horizon constraints to refine the estimated RGT field.

Figure 4a shows the result obtained with 2D stratigraphic constraints. In the direct inference result, RGT-Est recovers the main stratigraphic framework and produces RGT contours that generally follow the seismic reflections. However, local deviations can still be observed in structurally complex intervals, especially where the reflectors are weak, discontinuous, or laterally variable. The purple dashed curves denote the ground-truth horizons used for visual comparison, and the yellow boxes highlight representative regions where the direct inference result shows local mismatch. After incorporating the 2D stratigraphic constraints, the predicted RGT contours become better aligned with the reference horizons and more consistent with the seismic reflectors. The improvement is particularly evident in the highlighted regions, where the constrained result reduces local drift and better honors the target stratigraphic geometry.

Figure 4b further evaluates the effect of 3D stratigraphic constraints. The input consists of the seismic volume and sparse 3D horizon surfaces, as shown in Fig. 4b-1. Before incorporating these constraints, the direct inference result captures the overall stratigraphic trend, but local inconsistencies remain in regions with rapid lateral changes in reflector geometry. In the red-box

region, the estimated contours show noticeable mismatch with the stratigraphic surfaces and reduced lateral consistency. After introducing the 3D constraints, the RGT field becomes more consistent with the input horizon surfaces and shows improved spatial continuity across the volume. The constrained result better preserves the relative ordering between adjacent strata and reduces local distortions in the highlighted complex zone.

Overall, the experiment demonstrates that RGT-Est can flexibly support 2D or 3D horizon constraints as stratigraphic prompts. These constraints significantly improve the accuracy of RGT estimation, the alignment between RGT contours and seismic reflectors, and the volumetric consistency of the predicted stratigraphic field. This capability enables RGT-Est to operate not only as a fully automatic inference model but also as an interactive interpretation tool that can incorporate interpreter-provided stratigraphic priors for higher-precision RGT construction.

3.3 Generalization on Challenging Field Surveys

Figure 5 presents representative RGT estimation examples on several challenging field surveys from different geological settings. From left to right, each row shows the input seismic volume, the seismic volume overlaid with RGT contours, and the estimated RGT field. These examples include the Costa Rica survey, the Poseidon survey from Australia, two Netherlands surveys, and two field surveys from a region in China. The tested volumes cover diverse structural and depositional conditions, including strong deformation, steeply dipping reflectors, complex fault systems, unconformity-like terminations, intrusive or diapiric structures, and laterally variable stratigraphic patterns.

The first row (Fig. 5a) shows the Costa Rica survey, where the seismic image contains strongly deformed reflectors and laterally discontinuous events. Such structural complexity makes it difficult to construct a globally ordered RGT field, because local reflector discontinuities can easily lead to contour distortion or stratigraphic inconsistency. RGT-Est produces a continuous RGT volume that follows the main seismic events across the deformed zones. The overlaid contours remain generally consistent with the reflector geometry, indicating that the method can recover a coherent stratigraphic framework even when the seismic reflections are locally disrupted.

The second row (Fig. 5b) shows the Poseidon survey from Australia. This example is characterized by a prominent large-scale structural high and strongly varying reflector dips. The estimated RGT field captures the overall uplift geometry and maintains a smooth age progression from shallow to deep intervals. In the contour-overlaid view, the predicted RGT contours conform well to the inclined reflectors on both flanks of the structure, suggesting that RGT-Est can adapt to complex three-dimensional structural relief and strong lateral variation in stratigraphic geometry.

The third row (Fig. 5c) shows a Netherlands survey with laterally continuous but highly variable reflectors, local deformation, stratigraphic terminations, and multi-stage depositional units. In this area, different stratigraphic packages are interleaved and locally overprinted, producing complex reflector relationships and strong lateral variations in seismic texture. Even under such multi-phase stratigraphic interference, RGT-Est generates a stable and coherent RGT field. The estimated contours remain well aligned with the seismic reflections and preserve a consistent relative ordering across different stratigraphic units, indicating that the method can robustly handle complex depositional superposition and laterally varying geological architectures.




Figure 5. Representative RGT estimation results on challenging field surveys. From left to right, each row shows the input seismic volume, the estimated RGT field, and the horizons extracted from the estimated RGT field and overlaid on the seismic volume. The six rows correspond to the Costa Rica survey, the Poseidon survey in Australia, two Netherlands surveys, and two field surveys from a region in China. These examples cover complex geological settings, including strong deformation, steeply dipping reflectors, faulted structures, multi-stage stratigraphic units, and diapiric or intrusive structures. RGT-Est produces coherent RGT fields and contours that generally follow the seismic reflectors across these challenging surveys.

The fourth row (Fig. 5d) shows another Netherlands survey with dense faulting, strong lateral structural variation, and poorly defined reflection events. In this area, the seismic reflections are weak, discontinuous, and locally obscured by fault-related deformation, making it difficult to identify continuous reflectors and maintain consistent RGT ordering across fault-bounded blocks. Despite these ambiguous seismic responses, RGT-Est produces contours that generally follow the main reflector trends on both sides of the fault zones and preserve a coherent stratigraphic framework across the volume. This result indicates that the proposed method can robustly handle unclear reflection axes and fault-related discontinuities while maintaining a stable relative geologic time field in the surrounding strata.

The fifth and sixth rows (Fig. 5e,f) show two examples from a field survey in China. The fifth row contains a densely faulted structural setting, where steep faults cut through the stratigraphic sequence and introduce abrupt discontinuities in the seismic image. RGT-Est successfully represents these discontinuities while maintaining a consistent ordering of the surrounding strata. The contours follow the main reflectors on both sides of the faulted zones and avoid severe crossing or inversion artifacts. The sixth row presents a more challenging case with a pronounced diapiric or intrusive structure, where the surrounding strata are strongly uplifted and deformed. Despite the severe structural disturbance, RGT-Est recovers a geologically reasonable RGT field. The contour-overlaid result shows that the predicted RGT contours bend consistently around the central structure and remain aligned with the deformed reflectors.

Overall, these examples demonstrate the strong field-data generalization ability of RGT-Est across different surveys, geological settings, and data qualities. The method can produce coherent RGT fields not only in relatively continuous stratigraphic intervals but also in challenging regions with strong deformation, steep dips, fault-related discontinuities, multi-stage stratigraphic interference, and complex structural relief. The consistent alignment between the estimated RGT contours and seismic reflectors indicates that RGT-Est effectively preserves stratigraphic ordering and provides a reliable basis for horizon correlation and subsequent structural interpretation.

3.4 Feature Visualization

To further examine what RGT-Est learns from seismic data, we visualize the intermediate feature representations extracted from Stage 3 of the HRNet backbone. Specifically, principal component analysis is applied along the channel dimension of the feature maps, and the first three principal components are rendered as RGB volumes. This visualization provides an intuitive way to inspect whether the learned features correspond to geologically meaningful structures rather than only low-level amplitude textures.

Figure 6 shows representative examples from different field surveys. For each case, the seismic volume is shown together with the RGB visualization of the corresponding intermediate features. The PCA-based feature maps exhibit strong responses around key geological structures, including faults, unconformities, continuous horizons, slope bodies, and deformed stratigraphic intervals. In faulted regions, the feature responses highlight discontinuous zones and fault-related reflector offsets. Near unconformities and slope structures, the responses delineate major stratigraphic boundaries and lateral changes in depositional geometry. In relatively continuous intervals, the features remain organized along reflector-parallel patterns, indicating sensitivity to stratigraphic layering.




Figure 6. Visualization of intermediate features learned by RGT-Est. For each example, the seismic volume is shown together with the RGB visualization of the Stage 3 HRNet feature map after PCA along the channel dimension. The feature responses highlight geologically important structures, including faults, unconformities, horizons, slope bodies, and deformed stratigraphic intervals. This indicates that RGT-Est learns structure-aware and stratigraphy-aware representations that support coherent RGT estimation.

These observations suggest that the HRNet backbone in RGT-Est learns structure-aware and stratigraphy-aware representations from seismic data. The network does not merely encode local amplitude variations, but also captures geological elements that are critical for constructing a coherent RGT field. Such feature organization helps explain why RGT-Est can preserve reflector alignment, maintain stratigraphic ordering, and generalize to structurally complex field surveys.

4 Discussion

The results show that RGT-Est provides a stratigraphically consistent way to estimate relative geologic time from 3D seismic data. Rather than treating RGT estimation as a purely numerical regression problem, the proposed method represents the RGT field in a sinusoidal phase space that is more closely related to the periodic and ordered nature of stratigraphic layering. This design is important for solid Earth applications because an RGT volume is not only a mathematical scalar field, but also a compact geological representation of depositional order, structural deformation, and stratigraphic correlation. When the estimated contours follow seismic reflectors and maintain a stable relative ordering, they provide a physically meaningful basis for interpreting horizons, unconformities, faults, and deformed stratigraphic packages.

The comparison with DeepRGT[†] demonstrates the importance of this stratigraphy-aware formulation. Voxel-space regression can recover the first-order depth trend, but it often produces local smoothing or contour distortion where seismic reflections are weak, discontinuous, dipping, or structurally disturbed. Such errors are not merely visual artifacts. Because RGT is used to extract horizons and correlate stratigraphic surfaces, local inconsistencies may propagate into incorrect horizon tracking, distorted depositional geometries, or unrealistic structural models. By optimizing in the sinusoidal space and combining pointwise, perceptual, and adversarial constraints, RGT-Est produces contours that better honor seismic reflector geometry and preserve

more coherent stratigraphic ordering. This improvement is particularly relevant for structurally complex basins, where horizon continuity is commonly interrupted by faults, unconformities, slope deposits, and multi-stage deformation.

The field examples further indicate that RGT-Est can generalize across different geological settings. The Costa Rica example contains strong deformation and laterally discontinuous reflections, the Poseidon example contains pronounced structural relief and steeply varying reflector dips, the Netherlands examples include slope structures and multi-stage stratigraphic units, and the China examples include densely faulted intervals and diapiric or intrusive structures. These cases represent common but challenging situations in seismic interpretation, where local seismic attributes alone are often insufficient to construct a globally consistent stratigraphic framework. The ability of RGT-Est to produce coherent RGT fields across these surveys suggests that the method learns stratigraphic organization beyond survey-specific amplitude patterns. This property is essential for applying AI-based interpretation methods to field data from different basins, acquisition conditions, and tectono-stratigraphic environments.

The constraint-guided experiments show another practically important aspect of the method. In seismic interpretation, fully automatic RGT estimation is useful for rapid regional screening and preliminary structural analysis, but high-precision geological interpretation often requires expert control through picked horizons or stratigraphic markers. RGT-Est naturally supports this workflow by accepting sparse 2D or 3D horizon constraints as stratigraphic prompts. The 2D constraints improve local alignment on the constrained section, whereas the 3D constraints provide volumetric control and improve spatial consistency across the seismic volume. This capability links data-driven prediction with interpreter-guided geological modeling. It also reduces the gap between automatic AI inference and traditional interpretation practice, where human expertise remains essential for resolving ambiguous seismic responses.

From a geological perspective, the estimated RGT volumes can support several downstream solid Earth applications. First, they provide a dense stratigraphic coordinate system from which horizons can be extracted automatically and compared across faults, folds, and unconformities. Second, they can help characterize depositional architectures such as slope systems, clinoform-like geometries, and laterally varying stratigraphic packages. Third, they can serve as structural constraints for implicit geological modeling, reservoir characterization, and basin-scale stratigraphic analysis. Because the RGT field encodes relative ordering rather than only local reflector appearance, it provides a useful intermediate representation between seismic images and higher-level geological models.

Several limitations should also be noted. First, the reliability of RGT-Est still depends on the quality of seismic imaging. In regions where reflectors are severely obscured by noise, multiples, acquisition footprint, poor illumination, or strong velocity-related distortion, the predicted RGT field may remain uncertain. Second, although the sinusoidal representation improves stratigraphic consistency, the present framework does not explicitly impose hard geological constraints, such as strict monotonicity, non-intersection of isosurfaces, or fault-specific displacement relationships. Third, sparse horizon constraints can improve the result, but their effectiveness depends on the accuracy, distribution, and geological representativeness of the provided constraints. Future work could incorporate uncertainty estimation, fault-aware structural constraints, well markers, and physics-informed geological priors to further improve the reliability of RGT estimation in highly complex tectonic settings.

Overall, RGT-Est provides a practical framework for learning stratigraphically consistent RGT volumes from 3D seismic data. By combining a stratigraphy-aware sinusoidal representation with optional horizon guidance, the method supports both automatic and interpreter-guided workflows. This makes it useful not only for improving AI-based RGT estimation, but also for advancing seismic interpretation toward more consistent, reproducible, and geologically meaningful structural and stratigraphic modeling.

5 Conclusions

We developed RGT-Est, a deep learning framework for estimating relative geologic time from 3D seismic data. The method maps the predicted RGT field into a multi-frequency sinusoidal phase space and jointly uses pointwise, perceptual, and adversarial losses to improve stratigraphic consistency and geological plausibility. Experiments on multiple field surveys show that RGT-Est produces RGT contours that better follow seismic reflectors and preserve more coherent stratigraphic ordering than voxel-space regression. The method remains effective in challenging settings, including weak-reflection zones, slope structures, faulted intervals, multi-stage stratigraphic units, and diapiric or intrusive structures. RGT-Est also supports sparse 2D or 3D horizon constraints as stratigraphic prompts, which significantly improve horizon alignment and volumetric consistency. These results indicate that sinusoidal-space optimization provides an effective and flexible strategy for constructing geologically meaningful RGT volumes for horizon extraction, structural interpretation, and subsequent geological modeling.

Code and data availability. The source code of RGT-Est is openly available at <https://github.com/douyimin/RGT-Est>, and the associated trained model weights and seismic datasets used to reproduce the results in this work are archived on Zenodo at <https://doi.org/10.5281/zenodo.20118902> (Dou, 2026).

Author contributions. Y.D. and X.W. conceived the study and designed the methodology. Y.D. implemented the framework, conducted the experiments, and prepared the figures. X.W. supervised the project and provided guidance on the seismic interpretation aspects. H.G. and Z.B. contributed to the discussion of methodology, data analysis, and result validation. All authors contributed to writing and revising the manuscript.

Competing interests. The authors declare that they have no conflict of interest.

References

- Bader, S., Spikes, K., and Fomel, S.: Missing well-log data prediction using Bayesian approach in the relative-geologic time domain, in: SEG Technical Program Expanded Abstracts, pp. 804–808, 2018.
- Bi, Z., Wu, X., Geng, Z., and Li, H.: Deep relative geologic time: A deep learning method for simultaneously interpreting 3-D seismic horizons and faults, *Journal of Geophysical Research: Solid Earth*, 126, e2021JB021 882, 2021.
- de Bruin, G., Hemstra, N., and Pouwel, A.: Stratigraphic surfaces in the depositional and chronostratigraphic (Wheeler-transformed) domain, *The Leading Edge*, 26, 883–886, 2007.
- Dou, Y.: RGT-Est: Datasets and Trained Model Weights for Learning Stratigraphically Consistent Relative Geologic Time from 3D Seismic Data via Sinusoidal Mapping [Dataset/Software], Zenodo, <https://doi.org/10.5281/zenodo.20118902>, [Dataset/Software], 2026.
- Figueiredo, A. M., Gattass, M., and Szenberg, F.: Seismic horizon mapping across faults with growing neural gas, in: 10th International Congress of the Brazilian Geophysical Society, pp. 1476–1481, 2007.
- Fomel, S.: Applications of plane-wave destruction filters, *Geophysics*, 67, 1946–1960, 2002.
- Fomel, S.: Predictive painting of 3D seismic volumes, *Geophysics*, 75, A25–A30, 2010.
- Gao, H., Wu, X., and Liu, G.: ChannelSeg3D: Channel simulation and deep learning for channel interpretation in 3D seismic images, *Geophysics*, 86, IM73–IM83, 2021.
- Geng, Z., Wu, X., Shi, Y., and Fomel, S.: Deep learning for relative geologic time and seismic horizons, *Geophysics*, 85, WA87–WA100, <https://doi.org/10.1190/geo2019-0252.1>, 2020.
- Geng, Z., Zhao, Z., Shi, Y., Wu, X., Fomel, S., and Sen, M.: Deep learning for velocity model building with common-image gather volumes, *Geophysical Journal International*, 228, 1054–1070, 2021.
- Grose, L., Laurent, G., Aillères, L., Armit, R., Jessell, M., and Caumon, G.: Structural data constraints for implicit modeling of folds, *Journal of Structural Geology*, 104, 80–92, 2017.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A.: Image-to-image translation with conditional adversarial networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1125–1134, 2017.
- Jolicœur-Martineau, A.: The relativistic discriminator: a key element missing from standard GAN, arXiv preprint arXiv:1807.00734, 2018.
- Karimi, P. and Fomel, S.: Stratigraphic coordinates: A coordinate system tailored to seismic interpretation, *Geophysical Prospecting*, 63, 1246–1255, 2015.
- Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, in: International Conference on Learning Representations (ICLR), 2015.
- Li, J. and Huang, W.: Automatic extraction of seismic data horizon across faults, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–11, 2022.
- Liao, Z., Zhu, P., Zhang, H., Li, Z., Li, Z., and Ali, M.: A deep learning-based seismic horizon tracking method with uncertainty encoding and vertical constraint, *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1–13, 2024.
- Lim, J. H. and Ye, J. C.: Geometric GAN, in: arXiv preprint arXiv:1705.02894, 2017.
- Liu, N., Huo, J., Li, Z., Wu, H., Lou, Y., and Gao, J.: Seismic attributes aided horizon interpretation using an ensemble dense inception transformer network, *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1–10, 2024.
- Lomask, J., Guitton, A., Fomel, S., Claerbout, J., and Valenciano, A. A.: Flattening without picking, *Geophysics*, 71, P13–P20, 2006.

- Lou, Y., Zhang, B., Fang, H., Cao, D., Wang, K., and Huo, Z.: Simulating the procedure of manual seismic horizon picking, *Geophysics*, 86, O1–O12, 2021.
- Luo, S. and Hale, D.: Unfaulting and unfolding 3D seismic images, *Geophysics*, 78, O45–O56, 2013.
- Shi, Y., Wu, X., and Fomel, S.: SaltSeg: Automatic 3D salt segmentation using a deep convolutional neural network, *Interpretation*, 7, SE113–SE122, 2019.
- Stark, T. J.: Unwrapping instantaneous phase to generate a relative geologic time volume, in: 73rd Annual International Meeting, SEG, Expanded Abstracts, pp. 1707–1710, 2003.
- Stark, T. J.: Relative geologic time (age) volumes—Relating every seismic sample to a geologically reasonable horizon, *The Leading Edge*, 23, 928–932, 2004.
- Wang, F., Wu, X., and Wang, H.: Seismic horizon identification using semi-supervised learning with virtual adversarial training, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–11, 2022.
- Wu, H., Li, Z., and Liu, N.: Variable seismic waveforms representation: Weak-supervised learning based seismic horizon picking, *Journal of Petroleum Science and Engineering*, 214, 110412, 2022.
- Wu, X. and Fomel, S.: Least-squares horizons with local slopes and multigrid correlations, *Geophysics*, 83, IM29–IM40, 2018.
- Wu, X. and Hale, D.: Horizon volumes with interpreted constraints, *Geophysics*, 80, IM21–IM33, 2015.
- Wu, X. and Janson, X.: Directional structure tensors in estimating seismic structural and stratigraphic orientations, *Geophysical Journal International*, 210, 534–548, 2017.
- Wu, X. and Zhong, G.: Generating a relative geologic time volume by 3D graph-cut phase unwrapping method with horizon and unconformity constraints, *Geophysics*, 77, O21–O34, 2012.
- Wu, X., Luo, S., and Hale, D.: Moving faults while unfaulting 3D seismic images, *Geophysics*, 81, IM25–IM33, 2016.
- Wu, X., Liang, L., Shi, Y., and Fomel, S.: FaultSeg3D: Using synthetic datasets to train an end-to-end convolutional neural network for 3D seismic fault segmentation, *Geophysics*, 84, IM35–IM45, 2019.
- Wu, X., Geng, Z., Shi, Y., Pham, N., Fomel, S., and Caumon, G.: Building realistic structure models to train convolutional neural networks for seismic structural interpretation, *Geophysics*, 85, WA27–WA39, 2020.
- Yang, F. and Ma, J.: Deep-learning inversion: A next-generation seismic velocity model building method, *Geophysics*, 84, R583–R599, 2019.
- Yang, J., Wu, X., Bi, Z., and Geng, Z.: A multi-task learning method for relative geologic time, horizons, and faults with prior information and transformer, *IEEE Transactions on Geoscience and Remote Sensing*, 61, 5907–720, <https://doi.org/10.1109/TGRS.2023.3264593>, 2023.
- Yu, S., Ma, J., and Wang, W.: Deep learning for denoising, *arXiv preprint arXiv:1810.11614*, 2018.
- Zeng, H., Backus, M. M., Barrow, K. T., and Tyler, N.: Stratal slicing—Part 1: Realistic 3-D seismic model, *Geophysics*, 63, 502–513, 1998.
- Zeng, H., Zhu, X., Zhu, R., and Zhang, Q.: Guidelines for seismic sedimentologic study in non-marine postrift basins, *Petroleum Exploration and Development*, 39, 295–304, 2012.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O.: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.
- Zhao, Z. and Zhao, J.-g.: Seismic horizon tracking based on the TransUnet model, *Geophysics*, 90, IM1–IM13, 2025.