
An Intriguing Property of Geophysics Inversion

Yinan Feng¹ Yinpeng Chen² Shihang Feng¹ Peng Jin³ Zicheng Liu² Youzuo Lin¹

Abstract

Inversion techniques are widely used to reconstruct subsurface physical properties (e.g., velocity, conductivity) from surface-based geophysical measurements (e.g., seismic, electric/magnetic (EM) data). The problems are governed by partial differential equations (PDEs) like the wave or Maxwell’s equations. Solving geophysical inversion problems is challenging due to the ill-posedness and high computational cost. To alleviate those issues, recent studies leverage deep neural networks to learn the inversion mappings from measurements to the property directly.

In this paper, we show that such a mapping can be well modeled by a *very shallow* (but not wide) network with only five layers. This is achieved based on our new finding of an intriguing property: *a near-linear relationship between the input and output, after applying integral transform in high dimensional space*. In particular, when dealing with the inversion from seismic data to subsurface velocity governed by a wave equation, the integral results of velocity with Gaussian kernels are linearly correlated to the integral of seismic data with sine kernels. Furthermore, this property can be easily turned into a light-weight encoder-decoder network for inversion. The encoder contains the integration of seismic data and the linear transformation without need for fine-tuning. The decoder only consists of a single transformer block to reverse the integral of velocity.

Experiments show that this interesting property holds for two geophysics inversion problems over four different datasets. Compared to much deeper InversionNet (Wu & Lin, 2019), our method achieves comparable accuracy, but consumes significantly fewer parameters.

¹Earth and Environmental Sciences Division, Los Alamos National Laboratory, USA ²Microsoft Research, USA ³College of Information Sciences and Technology, The Pennsylvania State University, USA. Correspondence to: Youzuo Lin <ylin@lanl.gov>.

1. Introduction




Figure 1. Illustration of the near-linear relation property between geophysical measurements and properties after applying integral transform. $\{\Phi_n\}$ and $\{\Psi_m\}$ are two families of kernels for integral transforms (e.g., sine and Gaussian). Here, the full waveform inversion from seismic data to velocity map is used as an example.

Geophysics inversion techniques are commonly used to characterize site geology, stratigraphy, and rock quality. These techniques uncover subsurface layering and rock geomechanical properties (such as velocity, conductivity), which are crucial in subsurface applications such as subsurface energy exploration, carbon capture and sequestration, groundwater contamination and remediation, and earthquake early warning systems. Technically, these subsurface geophysical properties can be inferred from geophysical measurements (such as seismic, electromagnetic (EM)) acquired on the surface. Some underlying partial differential equations (PDEs) between measurements and geophysical property exist, where inversion gets its name. For example, velocity is reconstructed from seismic data based on full waveform inversion (FWI) of a wave equation, while conductivity is recovered from EM measurements based on EM inversion of Maxwell’s equations.

However, these inversion problems can be rather challenging to solve, as they are ill-posed. Recent works study them from two perspectives: *physics-driven* and *data-driven*. The former approaches search for the optimal geophysical property (e.g., velocity) from an initial guess, such that the generated geophysical simulations based on the forward modeling of the governing equation are closed to the real measurements (Virieux & Operto, 2009; Feng & Schuster, 2019; Feng et al., 2021). These methods are computationally expensive as they require iterative optimization per sample.

The latter methods (i.e., data-driven approaches) (Wu & Lin, 2019), inspired by the image-to-image translation task, employ encoder-decoder convolution neural networks (CNN) to learn the mapping between physical measurements and geophysical properties. *Deep* network architecture that involves multiple convolution blocks is employed as both encoder and decoder, which also results in heavy reliance on data and very high computational cost in training.

In this paper, we found an intriguing property of geophysics inversion that can significantly simplify data-driven methods as:

Geophysical measurements (e.g., seismic data) and geophysical property (e.g., velocity map) have near-linear relationship in high dimensional space after integral transform.

Let $u(x, t)$ denote a spatio-temporal geophysical measurement along horizontal x and time t dimensions, and $y(x, z)$ denote a 2D geophysical property along horizontal x and depth z . Since, in practice, geophysical measurement is mostly collected at the surface, and people want to invert the subsurface geophysical property, measurement u only contains spatial variable x , while property y includes (x, z) . As illustrated in Figure 1, the proposed property can be mathematically represented as follows:

$$\begin{aligned} \mathbf{U} &= [U_1, \dots, U_N]^T, \quad U_n = \iint u(x, t) \Phi_n(x, t) dx dt, \\ \mathbf{Y} &= [Y_1, \dots, Y_M]^T, \quad Y_m = \iint y(x, z) \Psi_m(x, z) dx dz, \\ \mathbf{Y} &\approx \mathbf{A} \mathbf{U}, \end{aligned} \quad (1)$$

where Φ_n and Ψ_m are kernels for integral transforms. After applying integral transforms, both geophysical measurement $u(x, t)$ and property $y(x, z)$ are projected into high dimensional space (denoted as \mathbf{U} and \mathbf{Y}), and they will have a near-linear relationship ($\mathbf{Y} \approx \mathbf{A} \mathbf{U}$). Note that the kernels ($\{\Phi_n\}, \{\Psi_m\}$) are *not* learnable, but well-known analytical kernels like sine, Fourier, or Gaussian.

Interestingly, this intriguing *property* can significantly simplify the encoder-decoder architecture in *data-driven* methods. The encoder only contains the integral with kernel $\{\Phi_n\}$ followed by a linear layer with weight matrix \mathbf{A} in Eq. 1. The decoder just uses a single transformer (Vaswani et al., 2017) block followed by a linear projection to reverse the integral with kernels $\{\Psi_m\}$. This results in a much shallower architecture. In addition, the encoder and decoder are learnt separately. The matrix \mathbf{A} in encoder can be directly solved by pseudo inverse and is frozen afterward. Only the transformer block and following linear layer in the decoder are learnt via SGD based optimizer.

Our method, named InvLINT (Inversion via LINear relationship between INTegrals), achieves comparable (or even

better) performance on two geophysics inversion problems (seismic full waveform inversion and electric/magnetic inversion) over four datasets (Kimberlina Leakage (Jordan & Wagoner, 2017), Marmousi (Feng et al., 2021), Salt (Yang & Ma, 2019), and Kimberlina-Reservoir (Alumbaugh et al., 2021)), but uses significantly less parameters than prior works. For instance, on Marmousi, our model only needs 1/20 parameters, compared to previous InversionNet.

2. Background

The governing equation of the seismic full waveform inversion is acoustic wave equation (Schuster, 2017),

$$\nabla^2 p(\mathbf{r}, t) - \frac{1}{c^2(\mathbf{r})} \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} = s(\mathbf{r}, t), \quad (2)$$

where $\mathbf{r} = (x, z)$ represents the spatial location in Cartesian coordinates (x is the horizontal direction and z is the depth), t denotes time, $c(\mathbf{r})$ is the velocity map, $p(\mathbf{r}, t)$ represents the pressure wavefield, ∇^2 is the Laplacian operator, and $s(\mathbf{r}, t)$ is the source term that specifies the location and time history of the source.

For the EM forward modeling, the governing equation is the Maxwell's Equations (Commer & Newman, 2008),

$$\begin{aligned} \sigma \mathbf{E} - \nabla \times \mathbf{H} &= -\mathbf{J}, \\ \nabla \times \mathbf{E} + i\omega \mu_0 \mathbf{H} &= -\mathbf{M}, \end{aligned} \quad (3)$$

where \mathbf{E} and \mathbf{H} are the electric and magnetic fields. \mathbf{J} and \mathbf{M} are the electric and magnetic sources. σ is the electrical conductivity and $\mu_0 = 4\pi \times 10^{-7} \Omega \cdot s/m$ is the magnetic permeability of free space.

3. Methodology

In this section, we use seismic full waveform inversion (from seismic data to velocity) as an example to illustrate our derivation of the linear property after integral transforms. We will also show the encoder-decoder architecture based on this linear property. Empirically, our solution is also applicable to EM inversion (from EM data to conductivity).

3.1. Near-Linear Relationship between Integral Transformations

In the following part, we will show the seismic data and velocity maps have the near-linear relation after integral transformation like the format of Equation 1. The seismic data p and velocity map c are governed by the wave equation (Equation 2). Note that seismic data p and velocity map c in wave equation corresponds to the input u and output y in Equation 1, respectively.

Rewriting wave equation in Fourier series: Similar to constant coefficients PDEs, we assume spatial variable

$\mathbf{r} = (x, z)$ and temporal variable t are separable, i.e., $p(x, z, t) = p_1(x, z)p_2(t)$, and $s(x, z, t) = s_1(x, z)s_2(t)$. Thus, Equation 2 is rewritten as

$$\begin{aligned} c^2(x, z)(\nabla^2 p_1(x, z)p_2(t) - s_1(x, z)s_2(t)) \\ = \frac{\partial^2}{\partial t^2}(p_1(x, z)p_2(t)). \end{aligned} \quad (4)$$

Next the temporal parts $p_2(t)$ and $s_2(t)$ are represented as Fourier series: $p_2(t) = \sum_{n=1}^N B_n e^{j2\pi nt}$ and $s_2(t) = \sum_{n=1}^N G_n e^{j2\pi nt}$. This turns Equation 4 as:

$$\begin{aligned} \sum_{n=1}^N c^2(x, z)(\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n)e^{j2\pi nt} \\ = \sum_{n=1}^N 4\pi^2 n^2 p_1(x, z)B_n e^{j2\pi nt}. \end{aligned} \quad (5)$$

To make sure both sides have the same coefficient for each n , the aggregation $\sum_{n=1}^N$ and $e^{j2\pi nt}$ can be removed from Equation 5 as:

$$\begin{aligned} c^2(x, z)(\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n) \\ = 4\pi^2 n^2 p_1(x, z)B_n. \end{aligned} \quad (6)$$

By further integrating over x , we have

$$\begin{aligned} \frac{1}{4\pi^2 n^2} \int c^2(x, z) |\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n| dx, \\ = \int p_1(x, z) |B_n| dx, \\ = \left| \iint \underbrace{p_1(x, z)p_2(t)}_{\text{Seismic data}} \underbrace{e^{-j2\pi nt}}_{\text{Fourier kernel}} dt dx \right|, \end{aligned} \quad (7)$$

where $|\cdot|$ is the modulus operator of complex numbers and $B_n = \int p_2(t)e^{-j2\pi nt} dt$ are the Fourier coefficients. Note that since B_n and G_n are complex numbers, we take module on both sides. Here, taking the real or imaginary part, rather than modulo, does not affect our conclusions. Now, the right hand of Equation 7 is in the same format with U_n in Equation 1. The kernel function $\Phi_n(x, t) = e^{-j2\pi nt} \mathbb{1}(x)$, where $\mathbb{1}(x) = 1$ for all x .

Approximation by Integral over z : In reality the seismic data $p(x, z, t)$ is mostly collected at the surface ($z = 0$). Thus, the right-hand side of Equation 7 is computable at $z = 0$. However, the left-hand side is hard to calculate, since $\nabla^2 p_1(x, z)$ and $s_1(x, z)$ are unknown. Here, we hypothesize that the left-hand side at $z = 0$ can be approximated by leveraging velocity map at multiple depth positions as:

$$\begin{aligned} \frac{1}{4\pi^2 n^2} \int c^2(x, z) |\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n| dx \Big|_{z=0} \\ \approx \iint c^2(x, z) F_n(x, z) dx dz, \end{aligned} \quad (8)$$

where $F_n(x, z)$ is the kernel function.

This hypothesis (Eq. 7–8) bridges integral transforms of the seismic data ($\iint p(t, x, z)e^{-j2\pi nt} dt dx|_{z=0}$) and velocity maps ($\iint c^2(x, z)F_n(x, z) dx dz$) via an auxiliary function $\frac{1}{4\pi^2 n^2} \int c^2(x, z) |\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n| dx|_{z=0}$.

It has two parts: (a) the double integral of velocity maps equals the auxiliary function, and (b) the 2D kernel $F_n(x, z)$ can be estimated by a set of basis functions, so we can further calculate the inverse problem we want to solve. The existence of $F_n(x, z)$ to achieve the former equality can be validated by a special case $F_n(x, z) = \frac{1}{4\pi^2 n^2} |\nabla^2 p_1(x, z)B_n - s_1(x, z)G_n| \delta(z)$ where $\delta(z)$ is an impulse function. The latter may weaken the former assertion of equality, but the misfit is likely small, as velocity map is continuous at most (x, z) positions and seismic data $p_1(x, z)$ and source $s_1(x, z)$ in the auxiliary function has strong correlation along x and z . Our experimental results over three datasets empirically validate this hypothesis.

Further simplification by a single kernel family: As discussed above, we simplify $F_n(x, z)$ as a weighted sum of a series of basis functions:

$$F_n(x, z) = \sum_{m=1}^M d_{n,m} \Psi_m(x, z), \quad (9)$$

where $d_{n,m}$ is the weight and $\Psi_m(x, z)$ is the basis function. By further plugging Equations 8 and 9 into Equation 7, we get

$$\begin{aligned} \sum_{m=1}^M d_{n,m} \iint c^2(x, z) \Psi_m(x, z) dx dz \\ \approx \left| \iint p(x, z, t) e^{-j2\pi nt} dt dx \right|_{z=0}. \end{aligned} \quad (10)$$

Relation to Equation 1: Equation 10 is special case of Equation 1. We can, therefore, express Equation 10 in the form of Equation 1 by letting:

$$\begin{aligned} y(x, z) &= c^2(x, z), & u(x, t) &= p(x, t), \\ A &= D^\dagger, & \Phi_n(x, t) &= e^{-j2\pi nt} \mathbb{1}(x), \end{aligned}$$

where A is pseudo inverse of matrix D and $D = [d_{n,m}]_{N \times M}$ is the matrix format.

In particular, $\mathbf{U} = [U_1, \dots, U_M]^T$ and $\mathbf{Y} = [Y_1, \dots, Y_N]^T$ are the high dimensional embeddings of the measurement




Figure 2. Schematic illustration of our proposed method, using seismic FWI as an example. The linear regression for two transformed embeddings is solved by pseudo inverse and is frozen afterward. The decoder is trained via SGD-based optimizer.

and geophysical property. $\{\Phi_n\}$ is chosen as cosine/sine or Fourier transform; while, based on the experiments, Gaussian kernel becomes our choice of the $\{\Psi_m\}$ to embed the spatial information in the geophysics property. It is true that the hypothesis may seem strong, however, its validity can be supported via our extensive experimental results using multiple datasets and various PDEs.

3.2. Simplified Encoder-Decoder Architecture

Based on the proposed mathematical property as shown in Equation 1, we can easily design a simple network architecture, accordingly. The encoder plays exactly the same role as the right-hand side of Equation 1, while the decoder, with a neural network, approximates the inverse mapping of the integral transformation ($\Psi_m^{-1}(x, z)$). The structure (Figure 2) of our InvLINT is described below.

Encoder: As illustrated in Figure 2, we design the encoder exactly the same to Equation 1, where an integral transformation with kernel $\{\Phi_n\}$, $n \in [1, N]$ is first implemented and followed by a linear layer represented by A . With such a simple linear relation, one can easily map the input measurement to the embedding of the output.

Decoder: There are many kernel functions (like Gaussian kernel), which does not have a close form inverse transformations. Instead, we use a shallow decoder network to approximate such a pseudo-inverse. To achieve this, we first use a linear layer L_1 to map Y to a more compact embedding and tile it a grid with the shape $R^{h \times w \times k}$. Here, h and w are the size of the velocity map with 32 times downsampling, and k is the number of channels. After that, $L_1(Y)$ is input into a 1-layer transformer, with patch size of $1 \times 1 \times k$. This shallow transformer is the only nonlinear part of our decoder.

The last parts of the model are a linear layer L_r . It upsamples each $1 \times 1 \times k$ patch to a $(32 + d) \times (32 + d)$ block¹, where d is a small integer. The final predicted velocity map \hat{c} can be construed by stitching all $h \times w$ blocks together. The

¹Since the size of output may not be divided exactly by 32, the recovered shape will be slightly different for different datasets.

purpose of this is to recover the output to the original shape with overlaps among blocks to remove the block effect.

3.3. Training

Because of the near-linear relation, we can easily solve the linear layer in the encoder, L_1 , with the least squares method. Specifically, we first compute the embedding of both encoder and decoder by integral transformations, calculate the solution of matrix A , and freeze it while training the decoder. The decoder is trained by an SGD-based optimizer. The loss function of our InvLINT is a pixel-wise MAE loss given as

$$\mathcal{L}(\hat{c}, c) = \ell_1(\hat{c}, c). \quad (11)$$

Peng et al. (Jin et al., 2022) find that combining MAE, MSE, and perceptual loss together is helpful to improve the performance. However, to make a fair comparison with the previous work, we only use MAE as our loss function.

4. Experiment

In this section, we present experimental results of our proposed InvLINT evaluated on four datasets and compare our method with the previous works, InversionNet (Wu & Lin, 2019) and VelocityGAN Zhang et al. (2019). We also discuss different factors that affect the performance of our method.

4.1. Implementation Details

4.1.1. DATASETS

In experiments, we verify our method on four datasets, of which three of them are used for seismic FWI, and one of which is for an EM inversion.

Kimberlina-Leakage: The geophysical properties were developed under DOE’s National Risk Assessment Program (NRAP). It contains 991 CO₂ leakage scenarios, each simulated over a duration of 200 years, with 20 leakage velocity maps provided (i.e., at every ten years) for each scenario (Jordan & Wagoner, 2017). Excluding some missing velocity maps, the data are split as 807/166 scenarios for training and testing, respectively. The size of the velocity maps is 401×141 grid points, and the grid size is 10 meters in both directions. To synthesize the seismic data, nine sources are evenly distributed along the top of the model, with depths of 5 m . The seismic traces are recorded by 101 receivers positioned at each grid with an interval of 15 m . The source frequency is 10 Hz. Each receiver collects 1251-timestep data for 1 second.

Marmousi: We apply the generating method in Jin et al. (2022), which follows Feng et al. (2021) and adopts the Marmousi velocity map as the style image to construct this low-

Dataset	Model	MAE↓	MSE↓	SSIM↑	#Parameters	FLOPs
Kimberlina Leakage	InversionNet (Wu & Lin, 2019)	9.43	1086.99	0.9868	15.81M	563.52M
	VelocityGAN Zhang et al. (2019)	9.73	1026.27	0.9863	16.99M	1.31G
	InvLINT (Ours)	8.13	1534.60	0.9812	1.49M (9.4%)	44.30M (7.9%)
Marmousi	InversionNet (Wu & Lin, 2019)	149.67	45936.23	0.7889	24.41M	189.58M
	VelocityGAN Zhang et al. (2019)	124.62	30644.31	0.8642	25.59M	259.49M
	InvLINT (Ours)	136.67	36003.43	0.7972	1.45M (5.9%)	9.31M (4.9%)
Salt	InversionNet (Wu & Lin, 2019)	25.98	8669.98	0.9764	13.74M	32.37M
	VelocityGAN Zhang et al. (2019)	332.62	145669.11	0.7760	14.92M	65.98M
	InvLINT (Ours)	24.60	8840.79	0.9742	1.62M (11.8%)	5.98M (18.5%)
Kimberlina-Reservoir	InversionNet (Wu & Lin, 2019)	0.01330	0.000855	0.9175	0.30M	1.20G
	VelocityGAN Zhang et al. (2019)	0.01313	0.000688	0.8611	1.48M	3.95G
	InvLINT (Ours)	0.00703	0.000537	0.9370	0.16M (53.3%)	96.10M (8.0%)

Table 1. Quantitative results evaluated on four datasets in terms of MAE, MSE and SSIM, the number of parameters and FLOPs. The percentages indicate the ratio of #Parameters (FLOPs) required by InvLINT to that required by InversionNet. Our InvLINT achieves comparable (or even better) inversion accuracy comparing to the InversionNet and VelocityGAN with a much smaller number of parameters and FLOPs.

resolution dataset. This dataset contains 30K with paired seismic data and velocity map. 24k samples are set as the training set, 3k samples are used as the validation set, and the rest are the testing set. The size of the velocity map is 70×70 , with the 10-meter grid size in both directions. The velocity ranges from $1,500m/s$ to $4,700m/s$. There are $S = 5$ sources placed evenly with a spacing of $170m$. The source frequency is 20 Hz. The seismic data are recorded by 70 receivers with a receiver interval of $10m$. Each receiver collects 1,000-timestep data for 1 second.

Salt: The dataset contains 140 velocity maps (Yang & Ma, 2019). We downsample it to 40×60 with a grid size of $10m$, and the splitting strategy 120/10/10 is applied. The velocity ranges from $1,500m/s$ to $4,500m/s$. There are also $S = 5$ sources used, with 12-Hz source frequency and a spacing of $150m$. The seismic data are recorded by 60 receivers with an interval of $10m$, too. Each receiver collects 600-timestep data for 1 second.

Kimberlina-Reservoir: The geophysical properties were also developed under DOE’s NRAP. It is based on a potential CO₂ storage site in the Southern San Joaquin Basin of California (Alumbaugh et al., 2021). We use this dataset to test our method in the EM inversion problem. In this data, there are 780 EM data as the geophysical measurement, and corresponding conductivity as the geophysical property. We use 750/30 as training and testing. EM data are simulated by finite-difference method (Commer & Newman, 2008) with two sources location at $x = 2.5km$, $z = 3.025km$ and $x = 4.5km$, $z = 2.5km$. There are $S = 8$ source frequencies from 0.1 to 8.0 Hz and recorded with its real and imaginary part. The conductivity is with the size of $351 \times 601 (H \times W)$, where the grid is $10m$ in all dimensions.

4.1.2. TRAINING DETAILS

The input seismic data and EM data are normalized to the range $[-1, 1]$. In practice, to supply more information, it always uses multiple sources to measure, where $s \in [1, \dots, S]$ is the index of different sources. After integration, all sources vectors will be concatenated. For the seismic data, we use $\Phi_n(x, t) = \sin(n\pi t)\mathbb{1}(x)/(x_{max} - x_{min})$ as the kernel function. However, for the EM data, since the raw data are already in the frequency domain and the input size is small, we skip the integral transformation step.

The Gaussian kernel can be represented as $\Psi_m(x, z) = \exp\left(-\frac{\|(x, z) - \mu_m\|_2^2}{2\sigma^2}\right)$. We let μ_m distribute evenly over the output shape. Then, the σ is set equal to the distance of adjacent μ . When applying Ridge regression to solve the linear layer in the encoder, and set the regularization parameter $\alpha = 1$.

We employ AdamW (Loshchilov & Hutter, 2018) optimizer with momentum parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$ and a weight decay of 1×10^{-4} to update decoder parameters of the network. The initial learning rate is set to be 1×10^{-3} , and we decay the learning rate with a cosine annealing (Loshchilov & Hutter, 2016), where $T_0 = 5$, $T_{mult} = 2$ and the minimum learning rate is set to be 1×10^{-3} . The size of every mini-batch is set to be 128. We implement our models in Pytorch and train them on 1 NVIDIA Tesla V100 GPU.

4.1.3. EVALUATION METRICS

We apply three metrics to evaluate the geophysical properties generated by our method: MAE, MSE and Structural Similarity (SSIM). In the existing literature (Wu & Lin, 2019; Zhang & Lin, 2020), both MAE and MSE have been employed to measure the pixel-wise error. SSIM is also con-

sidered to measure the perceptual similarity (Jin et al., 2022), since both velocity maps and conductivity have highly structured information, and degradation or distortion can be easily perceived by a human. Note that when calculating MAE and MSE, we denormalize geophysical properties to their original scale while we keep them in the normalized scale $[-1, 1]$ for SSIM according to the algorithm.

Moreover, we also employ two common metrics to measure the complexity and computational cost of the model: the number of parameters (#Parameters) and Floating-point operations per second (FLOPs).

4.2. Main Results

Table 1 shows the comparison results of our method with InversionNet on different datasets. Overall, our method achieves comparable or even better performance with a smaller amount of parameters and lower FLOPs. Below, we will provide in detail the comparison of all four datasets. It may be worthwhile mentioning that FWI is a quantitative inversion technique, meaning that it will yield both the shape and the quantitative values of the subsurface property.

Results on Kimberlina-Leakage: Compared to InversionNet and VelocityGAN, our method outperforms in MAE, slightly worse in MSE and SSIM. However, our InvLINT only needs less than 1/10 parameters and FLOPs. This demonstrates the power of our model, and further validates the properties we found. The velocity maps inverted by ours and InversionNet are shown in the first two rows of Figure 3. In the second example, despite of some noise produced by our method in the background, the CO₂ leakage plume (most important region as boxed out in green) has been very well imaged. Compared to ground truth, our method yields even better quantitative values than that obtained by InversionNet.

Results on Marmousi: Marmousi is a more challenging dataset due to its more complex structure. Compared to InversionNet, our method outperforms in all three metrics with significantly less computational and memory cost (about 1/20 parameters and FLOPs). This result again demonstrates not only the power of our model but also the validity of the near-linear relationship that we found. However, in such a large and complex dataset, VelocityGAN outperform others, where the GAN structure helps generating better results. The velocity maps inverted by ours and InversionNet are illustrated in the third and fourth rows of Figure 3. Our InvLINT and InversionNet perform comparably in both the shallow and deep regions compared to the ground truth.

Results on Salt: Compared to InversionNet, our method outperforms in MAE, and is slightly worse in MSE and SSIM with a very small gap. Moreover, our method uses 1/8 parameters and 1/5 FLOPs to those of InversionNet.

Note that, in this challenging dataset, which only has a small number of samples, VelocityGAN cannot converge well and yields bad results. This is a side effect of its complex structure. The velocity maps inverted by ours and InversionNet are illustrated in the fifth and sixth rows of Figure 3. Consistent with quantitative results, both methods generate similar results. In the shallow region, our method output a slightly clear structure; but in a deeper region (e.g., the red region in the first example), the output of InversionNet is a little close to the ground truth. However, the overall difference can be hard to distinguish. Our method achieves comparable results with much less complexity.

Results on Kimberlina-Reservoir: Compared to InversionNet and VelocityGAN, our method outperforms in all three metrics, with 1/2 parameters and 1/12 FLOPs to those of InversionNet. Because of the compact input, all model utilize the much smaller number of parameters. However, due to the simple architecture, InvLINT yields significantly fewer FLOPs but achieves better inversion accuracy. The conductivity results inverted by different models are shown in the last two rows of Figure 3. Contrary to previous results on the Kimberlina-Leakage dataset, our model yields clearer results. In the first example, we can see that the outputs of our model are less noisy; and in the second case, InvLINT inverts the deep region more precisely, as highlighted by the red squares. This is also consistent with the quantitative results.

At the same time, we find that the number of parameters of our model varies less for the same inverse problem. The number of model parameters is relatively independent of data size. In contrast, the previous methods are greatly affected by the input and output sizes. Moreover, our model not only requires fewer parameters, but also enables more *efficient* training and inference. When training on Marmousi dataset using 1 GPU (NVIDIA Quadro RTX 8000), our model is **9 times faster** than InversionNet/VelocityGAN (1 hour vs. 9 hours). We also tested inference runtime with batch size 1 on a single thread of an Intel(R) Xeon(R) CPU Gold 6248 v3 (2.5GHz). Our model is **16 times faster** than InversionNet/VelocityGAN (5 ms vs. 80ms). The small model size is suitable for memory-limited mobile devices. More visualization results are provided in the Appendix for readers who might be interested.

4.3. Ablation Tests

In this part, we will test how different kernel functions and network architectures will influence the performance of our method. We put our default setting at the first row of each table. For ease of illustration, we only provide results on Marmousi. Results on Kimberlina Leakage are given in the Appendix.

Different Encoder Kernels

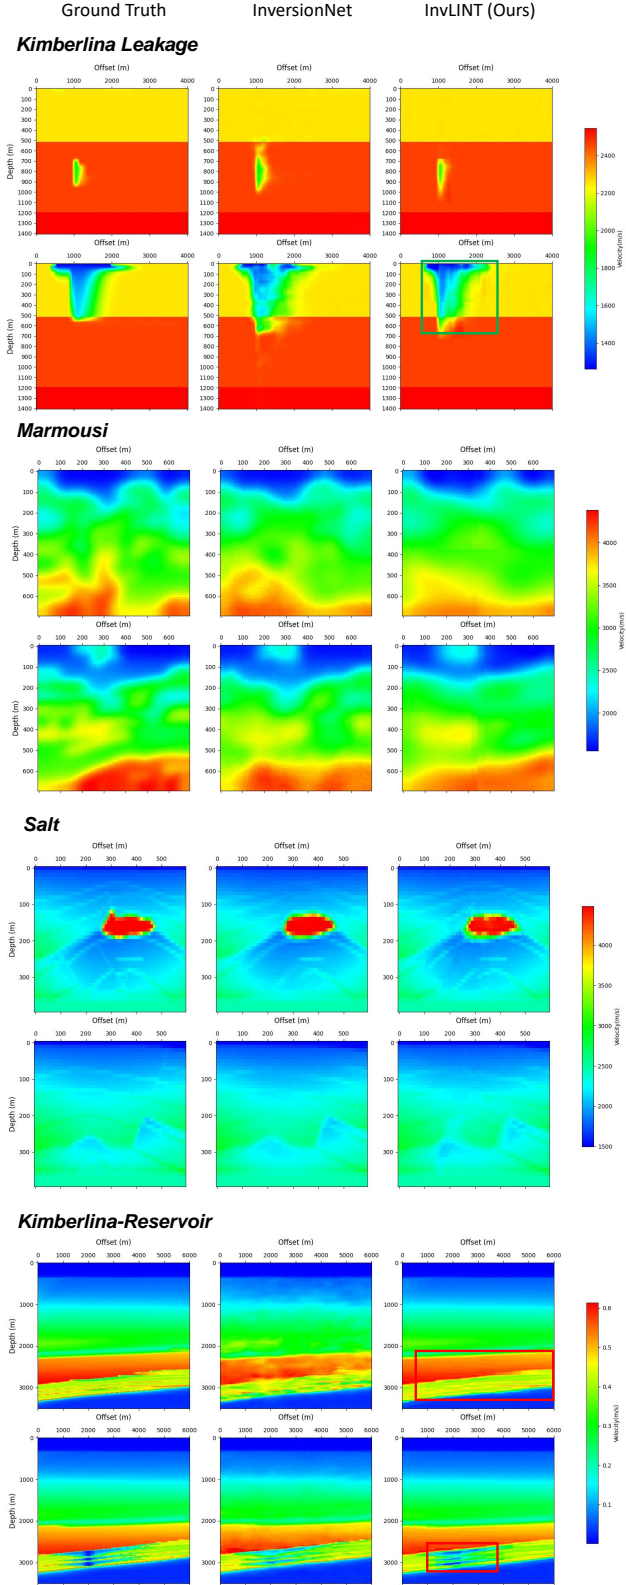


Figure 3. Illustration of results evaluated on four datasets

First, we conduct experiments by replacing the 1D sine kernel in the encoder to different 2D sine kernels. The quantitative results are shown in the Table 2. By comparing the results in Marmousi and the results in Kimberlina-Leakage (shown in the Appendix), we can see that the optimal strategy to integrate over x axis is distinct for different datasets. In Marmousi, using kernel $\sin(n\pi t) \cos(n\pi x)$ can improve the performance a lot. This kernel, however, does not perform well on other datasets (e.g., Kimberlina-Leakage).

Dataset	Encoder Kernel	MAE↓	MSE↓	SSIM↑
Marmousi	$\sin(n\pi t) \mathbb{1}(x)/(x_{max} - x_{min})$	136.67	36003.43	0.7972
	$\sin(n\pi t) \sin(n\pi x)$	138.76	37648.80	0.8042
	$\sin(n\pi t) \cos(n\pi x)$	128.33	32451.22	0.8115
	$\cos(n\pi t) \sin(n\pi x)$	140.14	38417.23	0.8031
	$\sin(n\pi(x+t))$	141.58	38383.58	0.7892
	$\sin(n\pi t) + \sin(n\pi x)$	142.12	38261.56	0.7884

Table 2. Quantitative results for Different Encoder Kernel.

Different Decoder Kernels

Then, we test different kernels for geophysical properties. In particular, we evaluate a series of 2D kernels: different 2D sine kernels, a sinc function kernel ($\sin(\pi \|\mathbf{r} - \mu_m\|_2) / \|\mathbf{r} - \mu_m\|_2$), and a Gaussian kernel with a smaller variance, noted as Gaussian σ . For the sinc function, the choice of μ_m is the same as the Gaussian kernel, while for Gaussian σ , we choose σ as 1/3 of the original. The quantitative results are shown in Table 3. As we can see, our choice of kernel outperforms the rest kernels. A smaller variance of Gaussian will yield a slightly worse result, while sinc kernel performs similarly to the Gaussian σ .


Dataset	Decoder Kernel	MAE↓	MSE↓	SSIM↑
Marmousi	Gaussian	136.67	36003.43	0.7972
	Sinc	138.02	36534.44	0.7952
	Gaussian σ	138.19	36579.46	0.7954
	$\sin(n\pi x) \sin(n\pi z)$	177.36	56102.75	0.7455
	$\cos(n\pi x) \sin(n\pi z)$	165.38	49463.79	0.7491
	$\sin(n\pi x) \cos(n\pi z)$	175.92	55424.26	0.7376
	$\sin(n\pi(x+z))$	209.47	74167.16	0.7057
	$\sin(n\pi x) + \sin(n\pi z)$	216.12	78496.77	0.7030

Table 3. Quantitative results for Different Decoder Kernel.

Different Number of Kernels

We also test different numbers of kernels for both Sine and Gaussian. We evaluate performance over a 6×6 grid where the dimensions of U and Y vary from 128 to 4096. The quantitative results are shown in Figure 4. Results indicate that the current selection of dimensions is appropriate. Obviously, reducing the model's size reduces its capacity, while choices of higher dimension are more prone to overfit. However, choosing a small dimension yields a smaller number of parameters and FLOPs. One can easily balance the performance and the cost based on his requirements and resources, indicating the flexibility of our model.

Different Decoder Architectures

Figure 4. Performance over dimensions of U and Y .

We aim to design an effective and efficient decoder to reverse the integral transform over a velocity map. The shifted Gaussian kernels used in integral transform split the velocity map into overlapping windows and encode the local structure within each window. To reconstruct the global structure of the velocity map from these local features, we leverage the transformer’s power in modeling long-range interaction in *a single layer*. Options like conv/deconv require more layers to cover long range.

To better illustrate this, we test the performance of different decoder architecture. Results are provided in Table 4. A transformer layer followed by a linear layer is a more accurate decoder than shallow conv/deconv layers. Deeper decoders with more conv/deconv layers achieve more accurate results, but require a larger model. When using the deconv decoder of InversionNet in our method, we achieve better performance, clearly outperforming InversionNet (MAE 126.6 vs. 149.7).

Dataset	Architecture	MAE \downarrow	MSE \downarrow	SSIM \uparrow	#Params	FLOPs
Marmousi	Transformer $\times 1$ + Linear*	136.67	36003.43	0.7972	1.45M	9.3M
	Conv $\times 2$ + Linear	140.72	37345.58	0.7903	0.30M	9.2M
	Deconv $\times 1$ + Linear	167.98	49728.14	0.7520	0.35M	10.8M
	(Deconv + Conv) $\times 5$	126.59	33830.73	0.8158	12.71M	94.6M
	(Up + Conv) $\times 5$	128.74	34854.78	0.8120	4.01M	56.7M

Table 4. Comparison among different decoder structures. (*) indicates the default decoder option.

Results for a Larger Decoder

Here, we evaluate our method with a larger/deeper decoder. Firstly, we test it using multiple unshared linear layers, rather than a shared one L_{r1} , in the last part of our decoder. Furthermore, we evaluate our model with a deeper transformer. The quantitative results are shown in Table 5. The result using unshared linear layers indicates that a single linear layer is enough and the model does not benefit from more parameters. On the other hand, a deeper transformer can improve the performance. Similar to the number of kernels, the balance is based on requirements.

Dataset	Architecture	MAE \downarrow	MSE \downarrow	SSIM \uparrow
Marmousi	1 layer Transformer	136.67	36003.43	0.7972
	Multi-Linear	138.82	36801.89	0.7939
	2 layers Transformer	134.24	35111.23	0.8002
	3 layers Transformer	132.19	34502.25	0.8037

Table 5. Quantitative results for a Larger Decoder.

4.4. Singular Value Analysis

Another major benefit of our simplified model is the ease of analysis. Since we use only one linear layer in the encoder, we can analyze it by performing singular value decomposition. The results are shown in Figure 5. Since the singular value varies greatly in different datasets, we divide it by its maximum value to normalize it and trunk it at 150 dim. Results indicate that for all datasets, the number of essential dimensions is less than 100. In other words, a 100-dimensional latent space is sufficient to represent the data. Specifically, we can see that a ten-dimensional latent space is enough for Kimberlina-Reservoir dataset. That answers why the required number of parameters of both our InvLINT and InversionNet on Kimberlina-Reservoir datasets are much smaller than that on other datasets. All in all, with such a simple architecture, our InvLINT is able to not only help in analyzing the problem but also help us quantify the difficulty of different datasets.




Figure 5. Normalized Singular Value Decomposition of the linear layer on different datasets.

4.5. Comparison to traditional FWI:

We performed new comparison with a widely used traditional FWI method (i.e., Multiscale FWI (Virieux & Operto, 2009)) on three seismic FWI datasets (Marmousi, Kimberlina-Leakage, Salt). Our method is consistently better on 3 datasets (MAE: 11.7 vs. 42.0 in Kimberlina-Leakage, 140.7 vs. 199.5 in Marmousi, 26.1 vs. 176.6 in Salt). The traditional FWI requires a good initial guess

and optimization per sample, resulting in slow processing (e.g., 4 hours per sample in Kimberlina-Leakage). Due to the limited rebuttal duration, we ran the comparison over 5 samples per dataset.

5. Related works

5.1. Data-driven Methods for FWI

Recently, based on deep learning, a new type of method has been developed. Araya-Polo et al. (2018) use a fully connected network to invert velocity maps in FWI. Wu & Lin (2019) consider the FWI as an image-to-image translation problem, and employ encoder-decoder CNN to solve. By using generative adversarial networks (GANs) and transfer learning, Zhang et al. (2019) achieved improved performance. In Zeng et al. (2021), authors present an efficient and scalable encoder-decoder network for 3D FWI. Feng et al. (2021) develop a multi-scale framework with two convolutional neural networks to reconstruct the low- and high-frequency components of velocity maps. A thorough review on deep learning for FWI can be found in Adler et al. (2021).

5.2. Physics-informed machine learning

Previous pure data-driven methods can be considered as incorporating physics information in training data. On the other hand, integrating the physics knowledge into loss function or network architecture is another direction. All of them are called Physics-informed neural networks (PINN). Raissi et al. proposed utilizing nonlinear PDEs in the loss function as a soft constraint (Raissi et al., 2019). Through a hard constraint projection, Chen et al. proposed a framework to ensure model's predictions strictly conform to physical mechanisms (Chen et al., 2021). Based on the universal approximation theorem of operators, in Lu et al. (2021), authors proposed DeepONet to learn continuous operators or complex systems. Sun et al. (2021) proposed a hybrid network design, which involves deterministic, physics-based modeling and data-driven deep learning. A comprehensive review of PINN can be found in Karniadakis et al. (2021).

6. Conclusion

In this paper, we find an intriguing property of geophysics inversion: *a near-linear relationship between the input and output, after applying integral transform in high dimensional space*. Furthermore, this property can be easily turned into a light-weight encoder-decoder network for inversion. The encoder contains the integration of seismic data and the linear transformation without fine-tuning. The decoder consists of a single transformer block to reverse the integral of velocity with Gaussian kernels.

Experiments show that this interesting property holds for two geophysics inversion problems over four different datasets. Compared to much deeper InversionNet, our method achieves comparable accuracy, but consumes significantly fewer parameters.

References

- Adler, A., Araya-Polo, M., and Poggio, T. Deep learning for seismic inverse problems: Toward the acceleration of geophysical analysis workflows. *IEEE Signal Processing Magazine*, 38(2):89–119, 2021.
- Alumbaugh, D., Commer, M., Crandall, D., Gasperikova, E., Feng, S., Harbert, W., Li, Y., Lin, Y., Manthila Samarasinge, S., and Yang, X. Development of a multi-scale synthetic data set for the testing of subsurface CO₂ storage monitoring strategies. In *American Geophysical Union (AGU)*, 2021.
- Araya-Polo, M., Jennings, J., Adler, A., and Dahlke, T. Deep-learning tomography. *The Leading Edge*, 37(1): 58–66, 2018.
- Chen, Y., Huang, D., Zhang, D., Zeng, J., Wang, N., Zhang, H., and Yan, J. Theory-guided hard constraint projection (hcp): A knowledge-based data-driven scientific machine learning method. *Journal of Computational Physics*, 445: 110624, 2021.
- Commer, M. and Newman, G. A. New advances in three-dimensional controlled-source electromagnetic inversion. *Geophysical Journal International*, 172(2):513–535, 2008.
- Feng, S. and Schuster, G. T. Transmission+ reflection anisotropic wave-equation travelttime and waveform inversion. *Geophysical Prospecting*, 67(2):423–442, 2019.
- Feng, S., Fu, L., Feng, Z., and Schuster, G. T. Multiscale phase inversion for vertical transverse isotropic media. *Geophysical Prospecting*, 69(8-9):1634–1649, 2021.
- Jin, P., Zhang, X., Chen, Y., Huang, S. X., Liu, Z., and Lin, Y. Unsupervised learning of full-waveform inversion: Connecting CNN and partial differential equation in a loop. In *Proceedings of the Tenth International Conference on Learning Representations (ICLR)*, 2022.
- Jordan, P. and Wagoner, J. Characterizing construction of existing wells to a co2 storage target: The kimberlina site, california. Technical report, National Energy Technology Laboratory (NETL), Pittsburgh, PA, Morgantown, WV , 2017.
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

- Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983, 2016.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In Sixth International Conference on Learning Representations (ICLR), 2018.
- Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. Nature Machine Intelligence, 3(3):218–229, 2021.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378:686–707, 2019.
- Schuster, G. T. Seismic inversion. Society of Exploration Geophysicists, 2017.
- Sun, J., Innanen, K. A., and Huang, C. Physics-guided deep learning for seismic inversion with hybrid training and uncertainty analysis. Geophysics, 86(3):R303–R317, 2021.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- Virieux, J. and Operto, S. An overview of full-waveform inversion in exploration geophysics. Geophysics, 74(6):WCC1–WCC26, 2009.
- Wu, Y. and Lin, Y. InversionNet: An efficient and accurate data-driven full waveform inversion. IEEE Transactions on Computational Imaging, 6:419–433, 2019.
- Yang, F. and Ma, J. Deep-learning inversion: A next-generation seismic velocity model building method. Geophysics, 84(4):R583–R599, 2019.
- Zeng, Q., Feng, S., Wohlberg, B., and Lin, Y. Inversion-net3d: Efficient and scalable learning for 3d full waveform inversion. arXiv preprint arXiv:2103.14158, 2021.
- Zhang, Z. and Lin, Y. Data-driven seismic waveform inversion: A study on the robustness and generalization. IEEE Transactions on Geoscience and Remote sensing, 58(10):6900–6913, 2020.
- Zhang, Z., Wu, Y., Zhou, Z., and Lin, Y. Velocitygan: Subsurface velocity image estimation using conditional adversarial networks. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 705–714. IEEE, 2019.

A. Appendix

A.1. Inversion Results of Different Datasets




Figure 6. Illustration of results evaluated on Kimberlina Leakage




Figure 7. Illustration of results evaluated on Marmousi

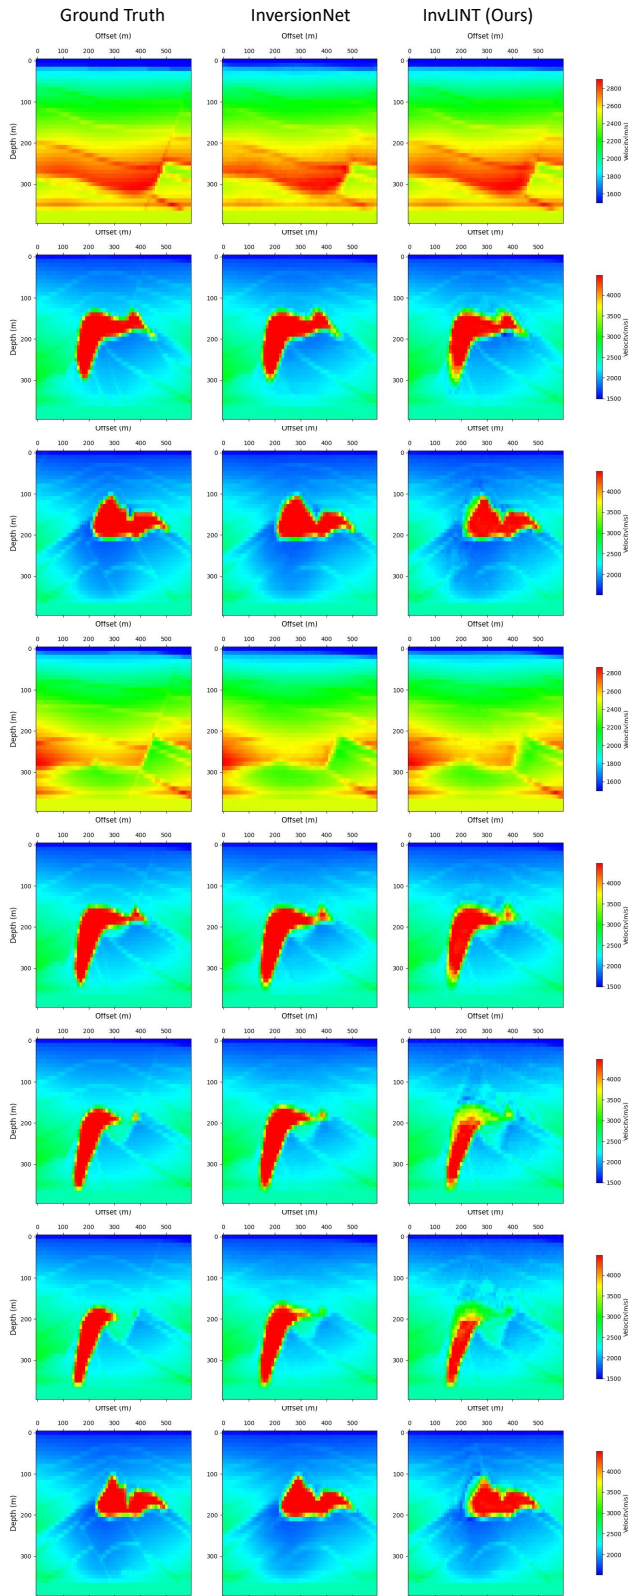


Figure 8. Illustration of results evaluated on Salt




Figure 9. Illustration of results evaluated on Kimberlina Reservoir

A.2. Ablation Test on Kimberlina Leakage

The ablation test results on Kimberlina Leakage are shown in Table 6-10

Dataset	Encoder Kernel	MAE↓	MSE↓	SSIM↑
Kimberlina Leakage	$\sin(n\pi t)\mathbb{1}(x)/(x_{max} - x_{min})$	8.13	1534.60	0.9812
	$\sin(n\pi t)\sin(n\pi x)$	11.07	3227.71	0.9783
	$\sin(n\pi t)\cos(n\pi x)$	8.88	2015.23	0.9804
	$\cos(n\pi t)\sin(n\pi x)$	10.95	3222.21	0.9782
	$\sin(n\pi(x+t))$	8.17	1751.89	0.9815
	$\sin(n\pi t) + \sin(n\pi x)$	8.10	1760.43	0.9817

Table 6. Quantitative results for Different Encoder Kernel.

Dataset	Decoder Kernel	MAE↓	MSE↓	SSIM↑
Kimberlina Leakage	Gaussian	8.13	1534.60	0.9812
	Sinc	8.90	2051.99	0.9789
	Gaussian σ	8.84	2042.94	0.9790
	$\sin(n\pi x)\sin(n\pi z)$	15.41	7357.48	0.9764
	$\cos(n\pi x)\sin(n\pi z)$	15.40	7349.02	0.9764
	$\sin(n\pi x)\cos(n\pi z)$	15.37	7252.45	0.9765
	$\sin(n\pi(x+z))$	12.86	4721.48	0.9767
	$\sin(n\pi x) + \sin(n\pi z)$	13.21	74719.95	0.9764

Table 7. Quantitative results for Different Decoder Kernel.

Dataset	#kernel	MAE↓	MSE↓	SSIM↑
Kimberlina Leakage	N=2048; M=512	8.13	1534.60	0.9812
	N=1024; M=512	8.63	1946.62	0.9811
	N=4096; M=512	8.29	1780.75	0.9808
	N=2048; M=128	8.76	2007.14	0.9805
	N=2048; M=1024	8.59	1898.95	0.9808

Table 8. Quantitative results for Different Number of Kernels.

Dataset	Architecture	MAE↓	MSE↓	SSIM↑
Kimberlina Leakage	Transformer $\times 1$ + Linear*	8.13	1534.60	0.9812
	Conv $\times 2$ + Linear	13.42	2447.81	0.9762
	Deconv $\times 1$ + Linear	21.32	4919.03	0.9648
	(Deconv + Conv) $\times 5$	6.86	1462.29	0.9841
	(Up + Conv) $\times 5$	6.87	1516.80	0.9840

Table 9. Quantitative results for Different Decoder Architectures. (*) indicates the default decoder option.

Dataset	Architecture	MAE↓	MSE↓	SSIM↑
Kimberlina Leakage	1 layer Transformer	8.13	1534.60	0.9812
	Multi-Linear	8.14	1799.70	0.9811
	2 layers Transformer	8.24	1781.50	0.9812
	3 layers Transformer	8.09	1707.53	0.9813

Table 10. Quantitative results for a Larger Decoder.

A.3. Regression Results for the Encoder Linear Layer

We also show here the regression results of the linear layer in our encoder on different datasets in Table 11. As a reference, we also show the range and mean of the regression target value as $y_{max}-y_{min}$, $|y_{mean}|$. The result demonstrate that how well the regressions are fitted.

Dataset	Set	MAE↓	MSE↓	$y_{max} - y_{min}$	$ y_{mean} $
Kimberlina Leakage	Training set	2.83	45.48	884.4	490.93
	Test set	4.63	245.09	885.27	492.65
Marmousi	Training set	4.26	33.38	107.05	4.92
	Test set	4.29	33.9	103.01	5.1
Salt	Training set	0.28	0.46	51.3	12.11
	Test set	0.48	1.98	49.35	11.97
Kimberlina Reservoir	Training set	169.25	2607.27	26497.1	7197.4873
	Test set	212.64	109288.08	26496.956	6849.95

Table 11. Quantitative results for a Larger Decoder.