Seismic compressive sensing by generative inpainting network: Toward an optimized acquisition survey

Xiaoyang Rebecca Li[1], Nikolaos Mitsakos[2], Ping Lu[2], Yuan Xiao[2], and Xing Zhao[3]

https://doi.org/10.1190/tle38120923.1

Abstract

The use of deep learning models as priors for compressive sensing tasks presents new potential for inexpensive seismic data acquisition. Conventional recovery usually suffers from undesired artifacts, such as oversmoothing, and high computational cost. Generative adversarial networks (GANs) offer promising alternative approaches that can improve quality and reveal finer details. An appropriately designed Wasserstein GAN trained on several historical surveys and capable of learning the statistical properties of the seismic wavelet's architecture is proposed. The efficiency and precision of this model at compressive sensing are validated in three steps. First, the existence of a sparse representation with different compression rates for seismic surveys is studied. Then, nonuniform samplings are studied using the proposed methodology. Finally, a recommendation is proposed for a nonuniform seismic survey grid based on the evaluation of reconstructed seismic images and metrics. The primary goal of the proposed deep learning model is to provide the foundations of an optimal design for seismic acquisition without a loss in imaging quality. Along these lines, a compressive sensing design of a nonuniform grid over an asset in Gulf of Mexico, versus a traditional seismic survey grid that collects data uniformly every few feet, is suggested, leveraging the proposed method.

Introduction

Compressive sensing (CS) (Candes et al., 2006; Donoho, 2006) provides methods that allow economic sampling from a sparse and compressible signal at acquisition time so that reconstruction of the data into dense regular intervals is plausible with sufficient confidence. As such, CS has significant potential to provide efficient seismic acquisition and imaging schemes. For example, when applied in the acquisition of seismic data in four spatial dimensions, CS can reduce the number of required geophone or hydrophone channels by a certain compression ratio toward each dimension under uniform or nonuniform sampling. Compared to a conventional survey design, CS offers significant advantages such as acquiring a survey faster and at lower cost and reconstructing the signals to an equivalent dense level later on, acquiring a much larger survey within similar timeline and budget, or designing extremely high-quality surveys at the same price.

We propose a framework that utilizes the power of generative adversarial networks (GANs) to model a manifold of seismic images from historical surveys. We use this GANs-based model to perform information recovery from sparse arrays.

Related works

The goal of image recovery is to reconstruct an image from its compressed peer. Two of the main application domains of image recovery are superresolution and inpainting. Superresolution aims at filling in the missing pixels in a low-resolution image, creating high-resolution uncompressed images, but it can only be combined with uniform sampling (Ledig et al., 2017; Wang et al., 2018). Pixel inpainting, on the other hand, can handle both uniform and nonuniform sampling scenarios by representing the nominal sampling scheme as arbitrary masks. In this work, we focus on the latter domain.

Conventional inpainting methods generally solve an interpolation problem using smooth function extensions, such as biharmonic interpolation (Damelin and Hoang, 2018) and Laplacian interpolation (Li et al., 2012). Both use regularized iterative reconstruction based on a predefined sparsity transform, which usually includes time-consuming iterative optimization and may result in undesired artifacts such as oversmoothing. Also, most of these methods suffer from image size limitations, high computational cost, and lack of refining details.

The use of deep generative models as priors for CS offers a promising alternative algorithmic approach for inexpensive seismic data acquisition. Deep generative models have been found to demonstrate superior performance at image retrieval tasks, improving quality and revealing finer details, compared to conventional approaches or pixel-wise deep learning models. The Wasserstein GAN (WGAN) (Arjovsky et al., 2017) is a well-known architecture that uses the Wasserstein distance to measure the distance between the data distribution learned by the generator and the actual data distrubution. The WGANs gradient penalty (WGANs-GP) loss function (Gulrajani et al., 2017) uses gradient penalty to additionally clip the weights of the discriminator. It succesfully addresses common problems observed in general GANs, such as unstable training caused by vanishing gradient or limited diversity of generated samples caused by mode collapse.

The generative inpainting network with contextual attention (GIN) (Yu et al., 2018) is a network that combines, and outperforms, several state-of-the-art approaches, including context encoders (Pathak et al., 2016), dilated convolutions of inpainting (Iizuka et al., 2017), and the aforementioned WGANs-GP. More importantly, this contextual-attention-based architecture, does not fuel its learning only from the known pixels surrounding the masked image area; it also looks for useful patches from other known image locations. For these reasons, we select the contextual-attention-based network as the basis for our seismic image compression method.

[1]University of Houston, Houston, Texas, USA. E-mail: xiaoyang.rebecca.li@gmail.com.
[2]Anadarko Petroleum Corporation, The Woodlands, Texas, USA. E-mail: nikolaos.mitsakos@anadarko.com; ping.lu@anadarko.com; yuan.xiao@anadarko.com.
[3]Texas A&M University, College Station, Texas, USA. E-mail: xingzhao@tamu.edu.

Although there have been attempts recently at applying deep learning methods to solve seismic image compression problems, such as superresolution by GAN (Lu et al., 2018; Picetti et al., 2018) or denoising and interpolation by autoencoder and reconstruction by GAN (Mandelli et al., 2019), none uses a pixel inpainting GAN applied on large amounts of real seismic data.

Goal of research

In 3D spaces, towed-streamer seismic data are usually acquired in subline directions (Figure 1). Two-dimensional sampling can be applied on both inline and crossline directions. The goal of this research is twofold. First, we use the state-of-the-art 2D generative inpainting network for sampled data reconstruction. Second, we explore the design of an optimal nonuniform sampling strategy, based on historical information. The following points describe our main contributions:

1) introduction of a GIN pixel inpainting model suitable for compressed seismic image recovery problems, under uniform or nonuniform sampling, capable of recovering the heavily sampled data efficiently and reliably;
2) superior model for CS (GIN-CS) on uniform sampling, that performs better than the original GIN and the state-of-the-art interpolation method for uniform sampling;
3) introduction of an effective nonuniform sampling survey recommendation, leveraging the GIN uniform sampling reconstructions and a hierarchical selection scheme.

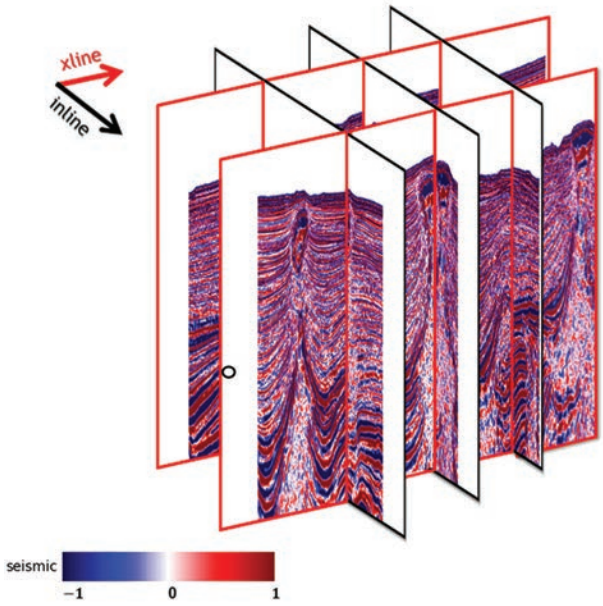

Figure 1. Compressive sampling in 3D space. 2D sampling can be applied on both inline and crossline images. The color bar shows the data range of the seismic image, normalized to –1 to 1.

As mentioned, we will apply our model for seismic image recovery when missing traces occupy a large portion of the image. In the following sections, we define the CS problem and describe the model design.

Sampling setting

Seismic image compression aims at efficient sparse representations of the signals. Provided a sampled signal and known sampling locations, a high-resolution image can be reconstructed with acceptable recovery loss. The sampling positions are indicated by binary masks (m) where 1 stands for missing, unknown pixels, and 0 stands for known pixels, as seen in Figure 2.

Compression rate (CR) is often used to describe the compression factor. It is defined as the proportion between uncompressed and compressed data size

$$\mathrm{CR} = \frac{\text{uncompressed data size}}{\text{compressed data size}}. \quad (1)$$

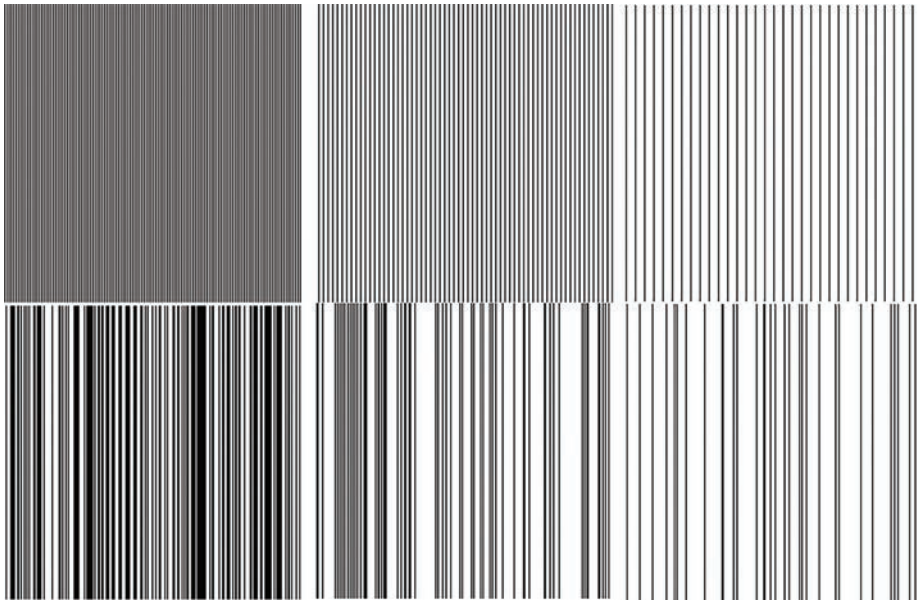

Figure 2. Examples of sampling binary masks for (a) uniform and (b) nonuniform sampling. For each of the 256 × 256 images, white areas indicate missing traces, and black areas indicate known traces. The compression rates, from left to right columns, are set to 2, 4, and 8, respectively.

CR is always larger than 1 when compressive sampling is applied. A higher CR value indicates more signal values are missing in the compressed image, thus potentially more challenging signal reconstruction. The CR value applies for both uniform and nonuniform sampling. As the bottom row in Figure 2 shows, nonuniform sampling may be regarded as the trace-wise shuffling of the missing traces arrangement in the uniform sampling under the same CR.

Bin width b defines the largest number of connected missing traces in a complete sampling survey, i.e., the maximum width of connected white areas in Figure 2. Bin width is crucial for pixel inpainting in nonuniform sampling because it indicates the largest gap that a model is required to fill. For a fixed value of b, the largest CR an

image could have is $W/(1+b)$, where W denotes the total number of traces.

Problem statement

The problem of compressed image recovery can be stated as a missing pixel inpainting problem. Given an incomplete image $\ddot{x}$, how do we fill the missing trace values corresponding to the binary mask m?

Using historical images of the uncompressed data set, we train a data-driven deep learning model utilizing the raw image (x) as ground truth and the binary mask to indicate the locations of the missing pixels. Recall that here we use 1 for unknown and 0 for known pixel locations. The predicted image output is denoted as $\tilde{x}$:

$$\tilde{x} = f(x, m). \tag{2}$$

Once the network has been trained, we can test the model's performance on any incomplete image $\ddot{x}$ from a different data set.

The main challenges in using an inpainting model to solve the seismic image sampling problems are:

1) Seismic images have significantly different statistical characteristics, such as texture-based patterns and a wide range of frequencies, compared to natural images.
2) The largest number of unknown pixels is only one-fourth of the full image size in the original inpainting network application, whereas the size of the unknown area in our task covers at least one-half of the image, in the CR = 2 case.
3) In our task, the known regions in the sampled image are sparsely distributed, contrary to the compact known regions addressed in the general inpainting problems.

We will address these problems by modifying the original network and employing it on different experiments.

Methodology

The GIN (Yu et al., 2018) is a feed-forward GAN with a contextual attention layer for missing pixel inpainting. Its architecture is composed by a coarse network and a refinement network, as seen in Figure 3. The coarse network fills in the missing pixels by dilated convolutions, as it only trains with the reconstruction loss to estimate a coarse pixel filling. The L_1 reconstruction loss is defined as

$$L_{reconst}(\tilde{x}) = \alpha_1 \left(\frac{\left\| (x - \tilde{x}) \odot (1 - m) \right\|_1}{\mathcal{N}_0 (1 - m)} \right). \tag{3}$$

Here, the masked image to be inpainted is described by $x \odot (1-m)$. $\mathcal{N}_0$ denotes the total number of zero values in the matrix, e.g., $\mathcal{N}_0(1-m)$ meaning number of unknown pixels, while α_1 is the L_1 loss penalty.

The refinement network applies contextual attention in parallel with dilated convolutions to enforce further detail consistency between the generated results and the ground truth images. Three strategies guide this refinement:

1) ***Decoupling global and foreground loss.*** The global loss term focuses at the differences between the whole ground truth and predicted images $(x, \tilde{x})$, whereas the foreground term only considers differences in the missing regions $\left(x \odot (1-m), \tilde{x} \odot (1-m)\right)$. The generator $G(x, \tilde{x})$ learns to create faithfully reconstruced images, guided by the differences between reconstructions and ground truth. At the same time, the discriminator $D(x, \tilde{x})$ constantly improves its ability to separate reconstructed images from images sampled directly from the data. Their loss functions are

$$L_G = L\left\{G(x, \tilde{x})\right\} + L\left\{G\left(x \odot (1-m), \tilde{x} \odot (1-m)\right)\right\} \tag{4}$$

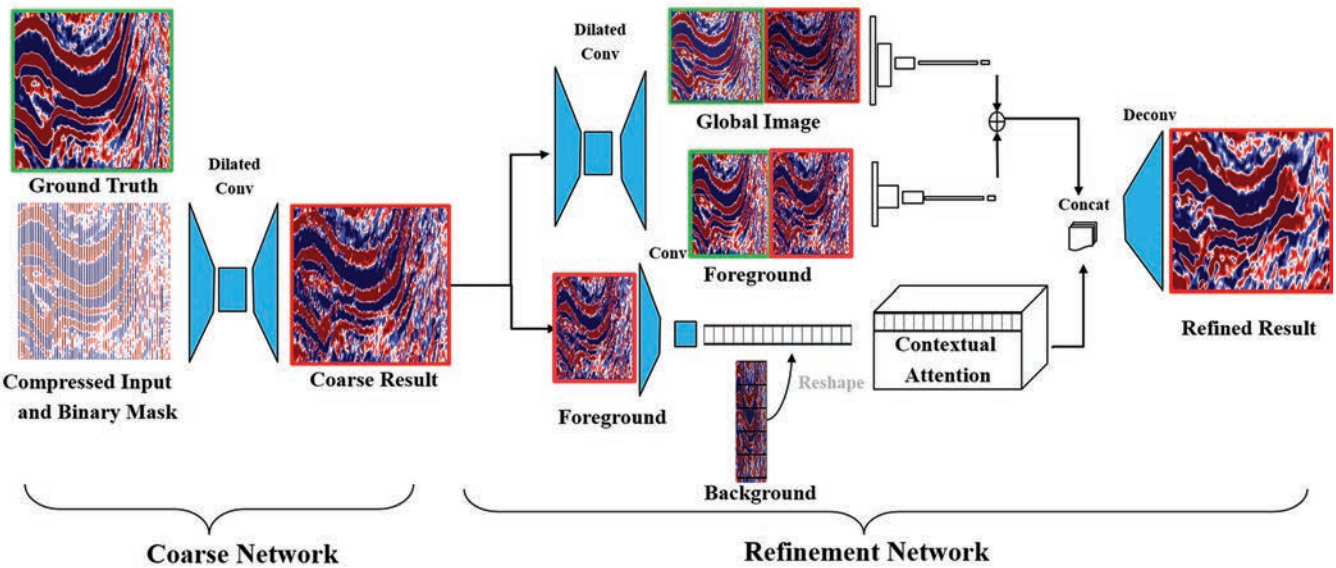

Figure 3. GIN architecture. A coarse network recovers the general features for the missing traces, followed by a refinement network that further reconstructs the finer structure. The ground truth and generated images are denoted by green and red frames, respectively. Foreground locations correspond to the unknown regions, while background locations are the known regions shaded as gray. Together, the two complete the whole image. (Seismic data courtesy of TGS)

and

$$L_D = L\{D(x,\tilde{x})\} + L\{D(x\odot(1-m),\tilde{x}\odot(1-m))\}. \quad (5)$$

2) ***Training with reconstruction loss concatenated with WGANs-GP adversarial loss.*** The WGANs-GP performs better than a general GIN in terms of vanishing gradient and mode collapse. The gradient penalty is defined as

$$P(x,\tilde{x}) = \left(\left\|\frac{\tilde{x}-\hat{x}}{\|\tilde{x}-\hat{x}\|_2}\right\|_2 - 1\right)^2, \quad (6)$$

where $\hat{x}$ correspond to random points between the straight lines connecting x and $\tilde{x}$. Since we use this to measure the distance between model distribution and read distribution, we replace the previous discriminator loss in equation 5 with the following term:

$$\begin{aligned} L_D &= L\{D(x,\tilde{x})\} + L\{D(x\odot(1-m),\tilde{x}\odot(1-m))\} \\ &+\lambda_{GP}L_{GP}\left(P(x,\tilde{x}) + P(x\odot(1-m),\tilde{x}\odot(1-m))\right) \end{aligned}. \quad (7)$$

3) ***Contextual attention layer.*** This layer aims at matching the foreground (unknown regions) and background (known regions) pixels. Compared to traditional convolutional neural networks, which are only able to learn the nearby patches at the same scale, the contextual attention attempts to learn the exact locations where background patches are used to generate unknown patches, no matter how far off. As Figure 3 illustrates, the background pixels are first extracted in 3 × 3 patches and then transformed into convolutional filters. The similarity of the foreground patch $f_{x,y}$ centered at pixel p_i location (x, y) and the background patch $b_{x',y'}$ centered at (x',y') is measured by the normalized inner product, scaled by channel-wise softmax

$$s_{x,y,x',y'} = softmax_{x',y'}\left(\lambda\left\langle\frac{f_{x,y}}{\|f_{x,y}\|},\frac{b_{x',y'}}{\|b_{x',y'}\|}\right\rangle\right), \quad (8)$$

where

$$softmax\,(p_i) = \frac{e^{p_i}}{\sum_{j=1}^{K} e^{p_j}} \quad \text{for } i = 1 \ldots K, \quad (9)$$

with K denoting the total number of pixels and λ being a constant. The final contextual attention layer is deconvoluted using background patches $b_{x',y'}$ as filters. The network can handle images with arbitrary resolutions since the contextual layer is fully convolutional and differentiable.

In our generative inpainting network for compressive sensing (GIN-CS), we replace the single bounding boxes used in the incomplete image generation with predefined binary masks. On one hand, a binary mask could be regarded as the combination of nonadjacent bounding boxes so that the multiple edges would lose the continuity of original spatial relations. On the other hand, the maximum width of connected missing traces we used is not larger than 16, so the edge effect would be ignored over the global image size. We should note that the original architecture uses spatially discounted reconstruction loss to reduce the edge effects in unknown pixel bounding boxes; however, due to the aforementioned difference in the shapes of our masks, we do not apply such a spatial discounting.

Measurement metrics

To assess the seismic image reconstruction performance, we use the mean squared error (MSE) and the peak signal-to-noise ratio (PSNR). For PSNR, a higher value implies better reconstruction quality, while the opposite holds for the MSE metric. We also adopt the structural similarity (SSIM) index (Wang et al., 2004) to describe the perceptual performance. The SSIM index varies from –1 to 1 with higher values implying that the structural similarity is better preserved.

Experimental setup

Data set description. The seismic data set we use in our experiments has been collected from the Gulf of Mexico (GoM), courtesy of TGS. It is characterized by different types of noise, such as constructional noise, and low-quality amplitude estimations under the huge salt bodies that cover most of the region, which makes recovery from sampled inputs challenging. We used a small portion of an internal offshore data set to train the network, and then the model was tested on the GoM data set. The two seismic data sets are characterized by substantially similar frequency content, intensities distribution, and geologic patterns. This allows us to explore the robustness of the generative inpainting models by training and testing on different data sets.

To analyze the seismic data in the form of an RGB image, we have clipped the original data values in the range of $[-\delta, \delta]$, where δ is the standard deviation of all the reflection amplitude values in the whole data set. Then, we rescale the values to the [–1, 1] range and apply the Matplotlib standard seismic color map as seen in Figure 1. These preprocessing steps aim only at improving the intensity contrast of the seismic images so that the feature patterns can be more effectively captured and learned by the convolutional filters.

Parameters. For training, we cropped 5000 of the processed offshore seismic images into 256 × 256 pixels and mixed the inline and crossline cases. The image size is selected based on the trade-off between working station storage capacity (16 GB) and training speed. Running 50 training epochs takes approximately 10 hours. The main parameters are shown in Table 1.

As mentioned in the Methodology section, there are two ways to arrange the training masks: random single bounding boxes, as in the original GIN, or predefined binary sampling masks, as in our modified GIN-CS. For testing, both methods use binary masks.

To investigate the influence of different training masks on both uniform and nonuniform sampling, we exploit six training mask types:

Table 1. Basic parameters for training.

L_1 loss α_1	GAN loss	GP loss λ_{GP}	Feature loss λ	Batch size	Iterations per epoch
1.2	0.001	2000	0.01	10	2000

- uniform sampling binary masks at four different compression rates, namely CR = 2, 4, 8, 16;
- random single bounding boxes with heights equal to image height (256 pixels) and largest bin width b = 8, 16 pixels.

Discussions on pretrain weights. We further investigate the influence of using pretrained GANs weights by training on the same set for 150 epochs using two different initialization schemes — one random and the other with weights pretrained on ImageNet (Deng et al., 2009). The PSNR/SSIM values we observed were 14.4 dB/0.33 for the ImageNet pretrained case and 25.8 db/0.91 for the random initialization. This indicates that transfer learning on seismic image data sets may not necessarily be beneficial to the overall training, contrary to what is often observed with natural images. This observation does not come as a complete surprise since seismic images share few similarities with the natural ImageNet images in terms of the statistical properties of their intensity distributions. So, when the number of seismic images in the training data set is much smaller than the number of samples in the ImageNet database, the network seems to have difficulty effectively adapting. At least, this may be implied by the observed performance differences in which a significant performance increase is evident when training from scratch, applying random weights. Further investigation is required to reach a clear determination on this, which is part of our ongoing research.

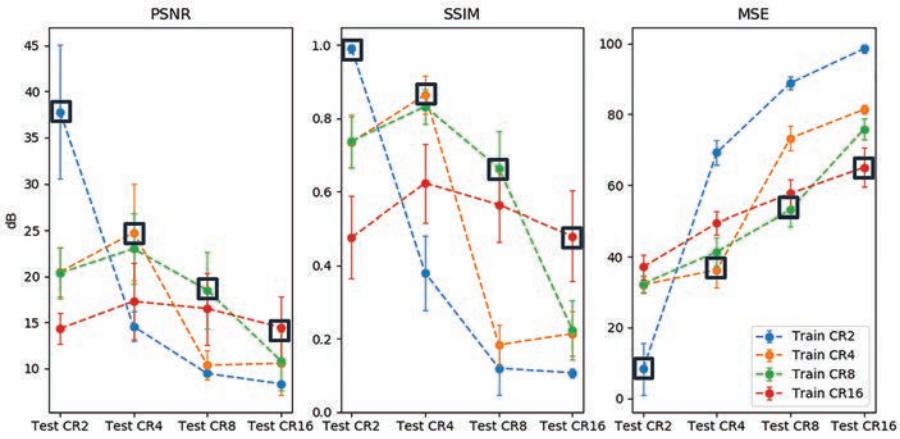

Figure 4. The error bar of (a) PSNR, (b) SSIM, and (c) MSE over 100 testing samples with different training and testing settings by GIN-CS. For each testing scenario, PSNR and SSIM reach the maximum values; MSE reaches the minimum values only when CS rates of training and testing are the same (squared dots).

Experiments on uniform sampling

Reconstruction by GIN-CS. First, we design an experiment to discover the relationship between CR in the training scheme and in testing samples. Four networks are independently trained using different CRs (CR = 2, 4, 8, 16) of the uniform binary GIN-CS masks. The evaluation results of these 16 scenarios on 100 images are shown in Figure 4. The charts imply that models trained under smaller CR (e.g., CR = 2) are more sensitive to testing CR as they show larger variations in the trace amplitudes. Furthermore, a higher CR training scheme does not seem to lead to better reconstruction when the testing CR is smaller. This implies that the surplus of known training information in the unknown testing region might be perceived as "noise" from the network. As Figure 4 shows, both PSNR and SSIM are maximized, while MSE is minimized, when training and testing CRs are equal. We conclude that the best setting for binary mask is CR (train) = CR (test).

Comparison with other methods. The original GIN model uses random single bounding boxes for training. By comparing the performance of our modified GIN-CS model with the same CR using binary mask networks (second column in Table 2) against the original GIN model with bounding boxes (third and fourth columns in Table 2), we can see our model improving the MSE and SSIM in all CR cases. Focusing on one trace from the testing image, we observe how our model's prediction aligns better with the ground truth (second column in Figure 5) relative to the GIN (third column in Figure 5).

Furthermore, we compare our GIN-CS with the conventional biharmonic method. As the first two columns in Figure 5 demonstrate, the GIN-CS models perform a more accurate signal prediction. Moreover, the evaluation matrices of MSE and SSIM are reported in Table 2.1 and 2.2, demonstrating the overall superior performance of our method.

Although for CR = 8, 16, our method does not achieve better performance in terms of PSNR (Table 2.3), it still generates closer-to-real seismic images without adding more artificial noise (GIN) or creating blurry fillings (biharmonic). An example of CR 8 is shown in Figure 6.

In terms of speed, the GIN-related methods are approximately 300 times faster than the traditional method, as seen in Table 2.4. More importantly, the run times are hardly influenced by the value of CR, contrary to the additional computational effort required by the biharmonic method.

Uniform sampling on whole data set. We test the GIN-CS on a 3D cube of the GoM data of approximate size 1500 × 1300 × 1000 pixels. There are 1500 crossline images of size 1300 × 1000 pixels each and 1300 inline images of 1500 × 1000 pixels.

Initially, we applied compression recovery on all the crossline images and inline images separately. Then, we combined the

Table 2. MSE, SSIM, PSNR, and runtime performance comparison of uniform sampling by different methods. GIN-CS* is our method. The highest values of each testing compression rates are bolded.

1) MSE	GIN-CS*	GIN (bbox8)	GIN (bbox16)	Biharmonic
CR2	***21***	32.6	32.4	27.5
CR4	***35.9***	43.7	46.8	51.3
CR8	***46.1***	52.6	59.7	68.4
CR16	***56.2***	64.5	60.7	81.9
2) SSIM	**GIN-CS***	**GIN (bbox8)**	**GIN (bbox16)**	**Biharmonic**
CR2	***0.9***	0.61	0.65	0.84
CR4	***0.7***	0.52	0.55	0.59
CR8	***0.44***	0.32	0.27	0.34
CR16	***0.22***	0.14	0.17	0.19
3) PSNR	**GIN-CS***	**GIN (bbox8)**	**GIN (bbox16)**	**Biharmonic**
CR2	***19.5 dB***	13.8 dB	14.7 dB	18.1 dB
CR4	***14.1 dB***	12.4 dB	13.1 dB	13.9 dB
CR8	10.4 dB	10.4 dB	9.70 dB	***11.2 dB***
CR16	8.0 dB	7.7 dB	8.2 dB	***9.6 dB***
4) Runtime(s)	**GIN-CS***	**GIN (bbox8)**	**GIN (bbox16)**	**Biharmonic**
CR2	2.9	***2.1***	2.1	162.0
CR4	***1.7***	1.9	1.7	354.7
CR8	***1.8***	1.9	1.8	465.8
CR16	2.0	1.7	***1.6***	506.2

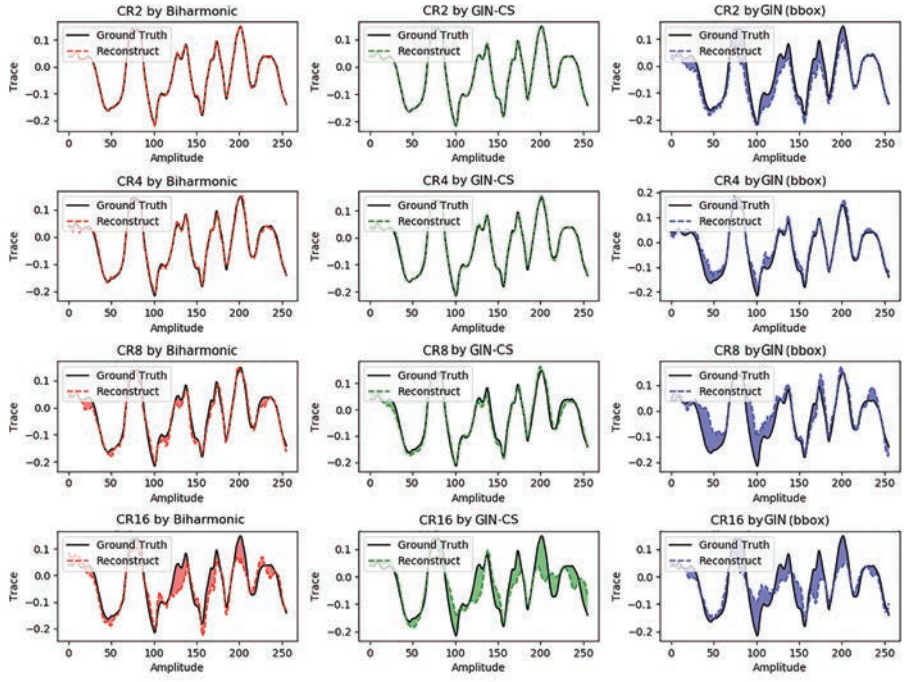

Figure 5. Trace comparison of uniform sampling with (a) biharmonic method (traditional sampling method), (b) GIN-CS (our method), and (c) GIN by bounding boxes (original GIN method). The smaller the shading areas are, the better the reconstruction performance is. GIN-CS performance is similar to the biharmonic method and a significant improvement over GIN.

results of crossline and inline by taking the average of the overlapped regions (depth view at depth = 500 is shown in Figure 7) in order to mimic the actual sampling in both dimensions at the same time. We are clear of the fact we only run the reconstruction on crossline and inline, and we plot it in depth only for validation purpose.

Experiments on nonuniform sampling

Random sampling. A nonuniform sampling scheme aims at recovering missing information under arbitrarily arranged, nonuniformly sized sampling scenarios only by restricting the maximum number of subsequent missing traces. To demonstrate the capability of the nonuniform sampling, we test a GIN model trained with bin width = 8 on several random sampling cases. The random sampling is independently generated by periodically shuffling the uniform sampling trace location. In this way, the nonuniform sampling series is able to imitate the real nonuniform acquisition scenarios. We report our findings in the last column of Table 4. We note that, in nonuniform cases, we examine the CR = 2, 3, 4 due to the restriction of bin width and CR.

Sampling survey recommendation. Thinking toward a robust construction scheme for a nonuniform optimal sampling survey setup, we also propose a sampling recommendation approach that leverages the fast implementation of image reconstruction with GIN. Our proposal is a fast, nonuniform sampling recommendation method, based on hierarchical uniform sampling, which requires only a small number of sampling test cases. We should note here that our recommended sampling method does not consider the connected sampling crossing section width, and its speed relies highly on the performance of the GIN.

Step 1: Mask generation. For a given uncompressed seismic image x of height H and width (total number of traces) W, we designed a set of binary masks

$$S = \{ m_i^{(b)} \mid b \in \{1,2,4,8\}, i \in \left\{0,1,2,3 \ldots \frac{B}{b} - 1\right\}\}$$

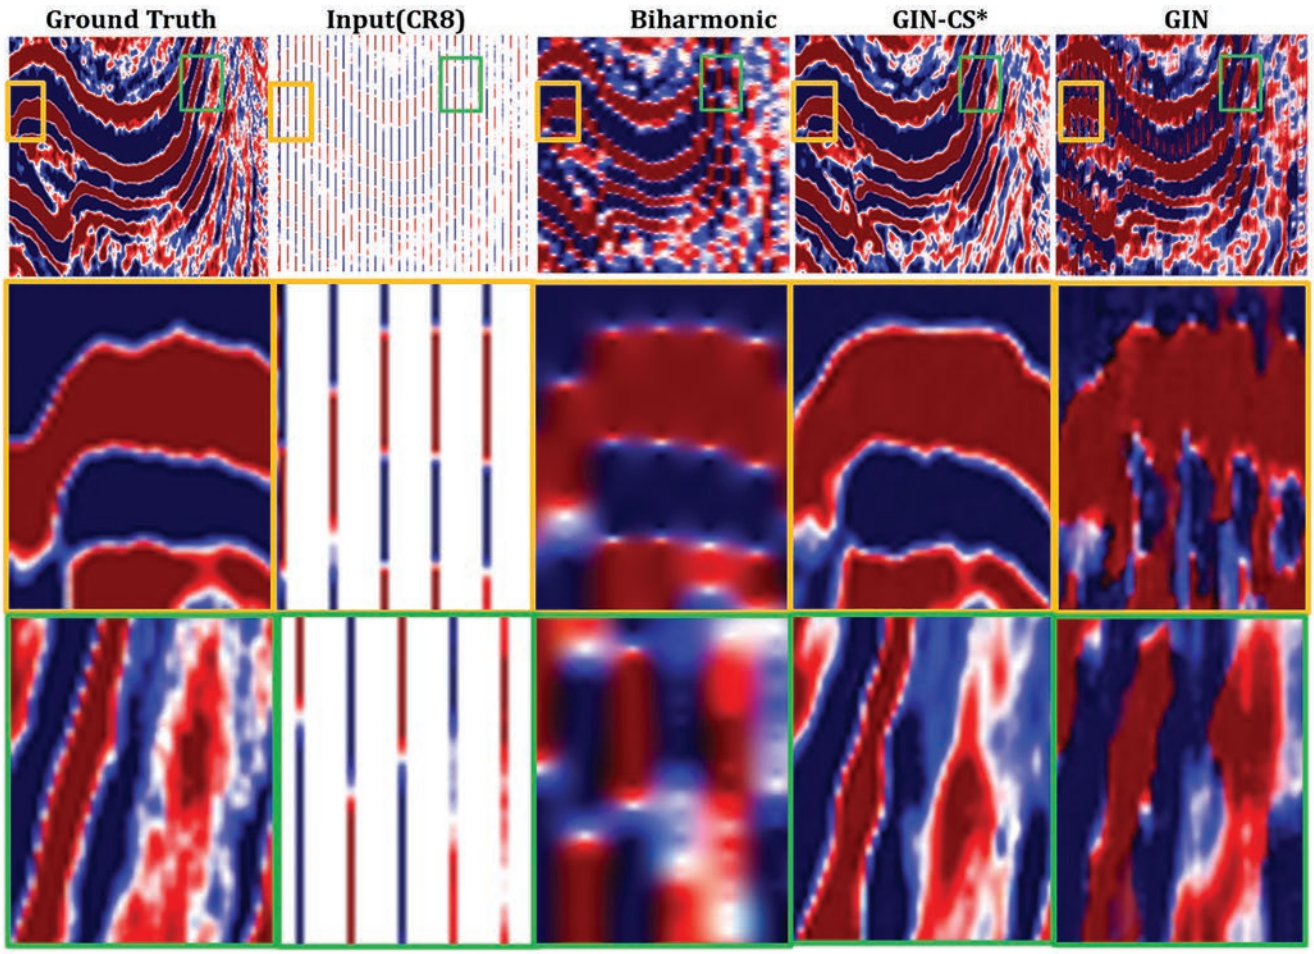

Figure 6. Reconstruction comparisons of CR 8. The second and third columns show closeup regions of CR 8 figures framed in yellow and green. (Seismic data courtesy of TGS)

Table 3. Example of a binary mask, $\text{mask}_1^{(4)}$. Unknown pixel is denoted as 1 and shaded in the table.

Trace	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
i	0				1				2				3				0				1			

where b is defined as the bin width of each group of connected missing traces and i is defined as the shifting index. In our experiments, we define section width $B = 2 \times b_{max} = 16$. In total, there are 16/1 + 16/2 + 16/4 + 16/8 = 30 masks in set S. An example is shown in Table 3.

Property 1. Within the subsets with same bin width $\{m_i^{(b)}\}$, all the masks are nonoverlapping, and together they cover the image completely:

$$\sum_{i=0}^{\frac{16}{b}-1} m_i^{(b)} = J_{H\times W} \ . \tag{10}$$

Step 2: Difference map generation. We test the image x with all the designed compression cases and generate 30 reconstruction results:

$$\overset{\text{X}}{\left\{m_0^{(1)}, m_1^{(1)}, \ldots m_1^{(16)}\right\}} \rightarrow \text{GIN} \rightarrow \left\{\tilde{x}_0^{(1)}, \tilde{x}_1^{(1)}, \ldots \tilde{x}_1^{(16)}\right\}$$

Then, we create the corresponding error matrix $\{Err_1^{(1)}, Err_2^{(1)}, \ldots, Err_2^{(16)}\}$ by calculating the pixel-wise square error of the reconstructions compared with ground truth.

$$Err_i^{(b)} = \left(x - \tilde{x}_i^{(b)}\right)^2_{H\times W} \ . \tag{11}$$

From Property 1, we know that the nonzero pixels in the error matrixes in each subset $\{Err_i^{(B)}\}$ are complementary to each other. So, summing them up, we form a complete image difference map and calculate its trace-wise mean vector with the size of $1 \times W$. For sampling by bin width b, the trace-wise average of difference map is defined as

$$Diff^{(b)} = \frac{1}{H}\sum_{b=1}^{H}\sum_{i=1}^{16/b} Err_i^{(b)} = \left| d_1^b, \ d_2^b, \ldots, \ d_w^b \right|_{1\times W}. \tag{12}$$

Step 3: Initial candidate traces generation. The difference map means vector is split into the individual difference values d_w for each

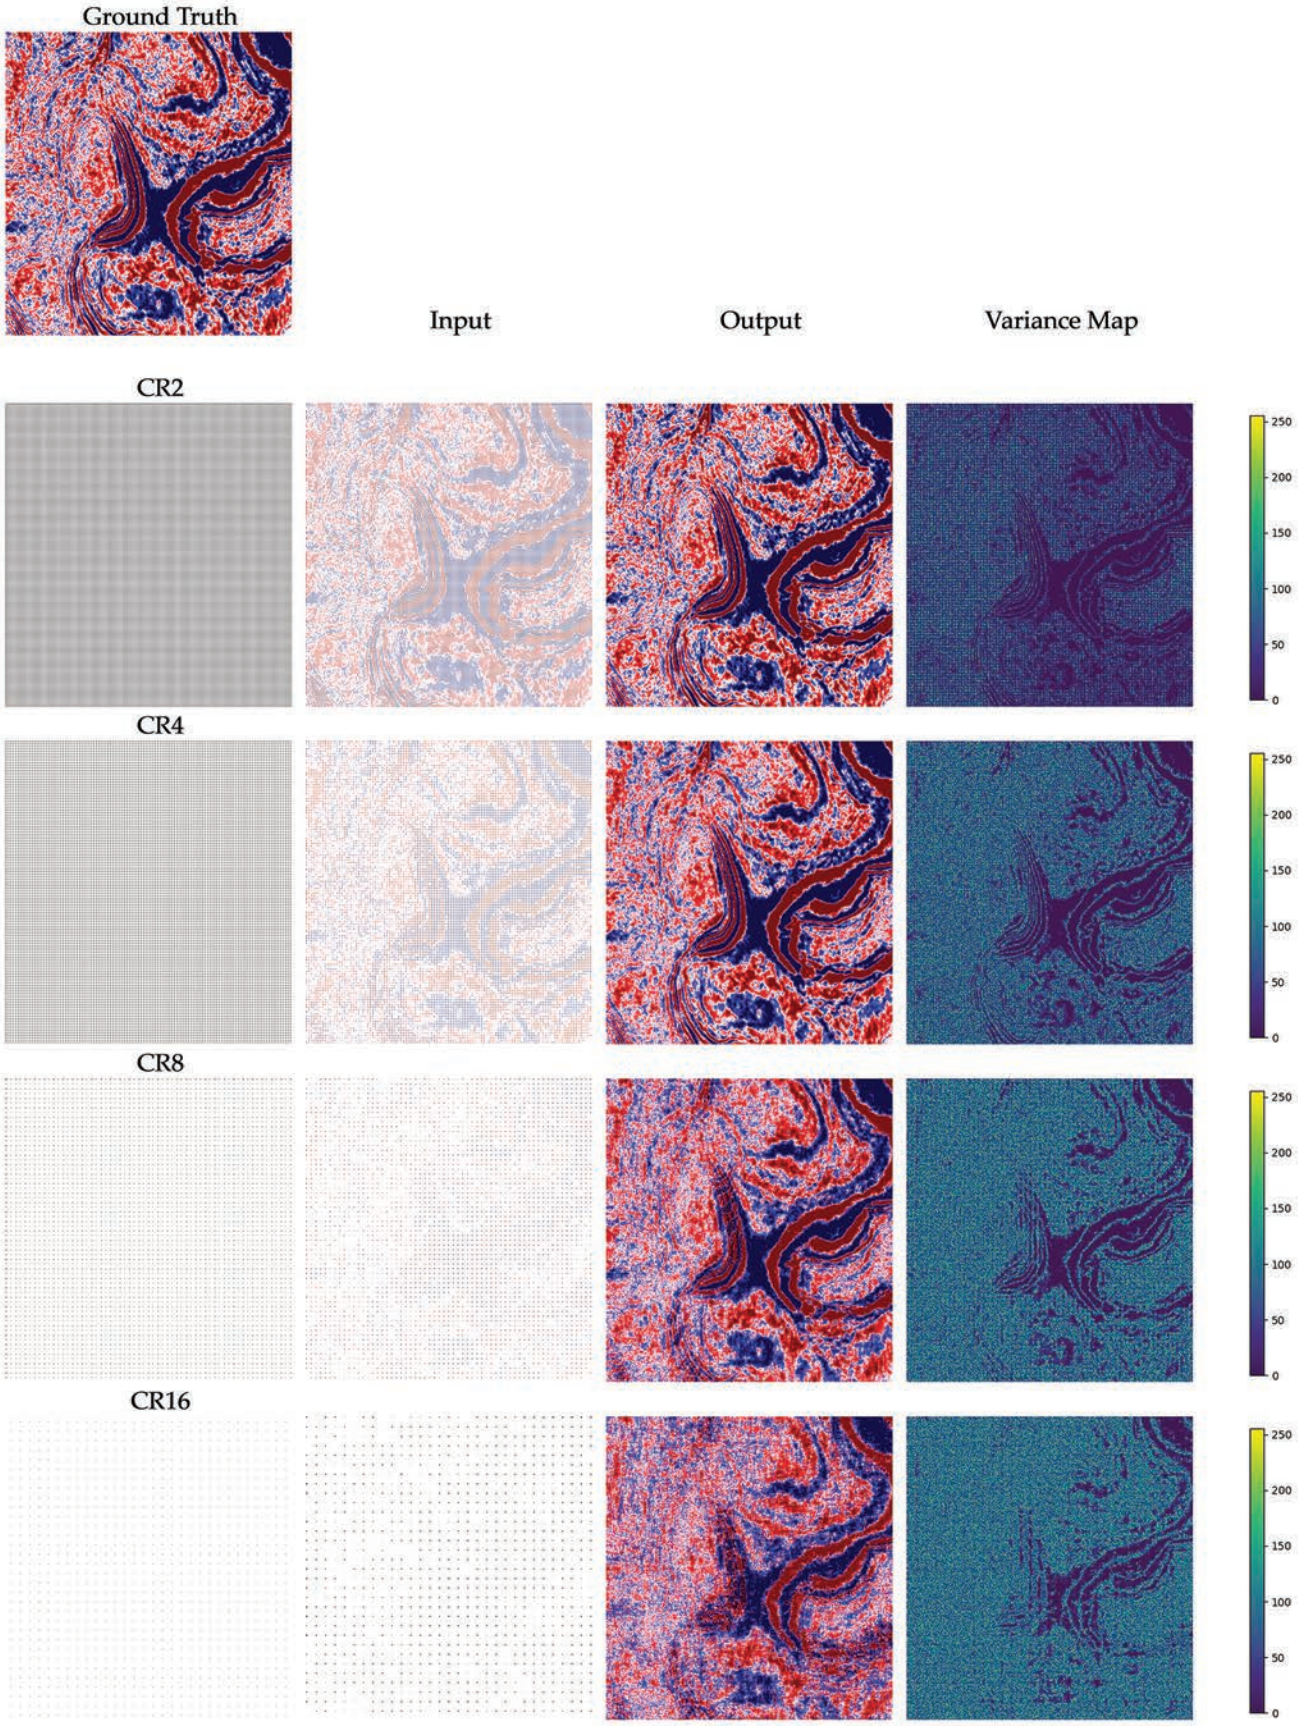

Figure 7. Stacked depth view comparing uniform sampling over CR = 2, 4, 8, 16 on a 512 × 512 cropped region of the GoM data set. All the crossline and inline sampled images are stacked together to form a 3D reconstruction of the whole block, showing its 2D depth view at depth = 500 only for validation. (Seismic data courtesy of TGS)

trace w. A smaller d_W^b value indicates better reconstruction at trace W. A plot example of the trace-wise average over bin width $b = 2, 4, 8$ for 64 traces is shown in the top left of Figure 8. (The $b = 1$ case is not displayed because it lies far below the other three curves.)

To compare d_w^b for each trace w without breaking the unknown connected traces, we introduce the observed interval $v \in \{2, 4, 8\}$ to distinguish from the concept of bin width b. The step mean difference is then defined by simply replacing the actual difference value d_w as its mean value over every v interval

$$sd_w^{b,v} = \frac{1}{v}\sum_{i=v\left[\frac{w}{v}\right]}^{v\left[\frac{w}{v}\right]+v} d_w^b \,. \tag{13}$$

For an observed interval v, the initial recommended candidate traces w' are the first trace indexes that reach the smallest step mean difference, when $b = v$.

$$w'^{(v)} = argmin_{w,v}\left(sd_w^{b,v}\right) \wedge (b = v)\ ,\ b \in \{2,4,8\}\,. \tag{14}$$

All candidate traces over the interval v form $cand^{(v)} = \{w'^{(v)}\}$, regarded as a subset of all traces. Figure 8 highlights the initial recommendation of bin width candidates $w'^{(v)}$ as solid dots. Of course, one trace might be a candidate for selection from several intervals or might never become a candidate for selection.

Step 4: Top-to-bottom hierarchical sorting.

Algorithm 1 Top-to-bottom hierarchical sorting

```
Total miss trace = 0
while Total miss trace < W(1 - 1/CR)  do
  Final recommended traces W*= { }
  Construct binary searching tree T̄
  for v = 8,4,2,1 do
       Sort ( cand (v) ) according to sd_w^{b,v}, where b = v
       for w'(v) in  cand(v) and w'(v) in T̄:
         W* = W* ∪ {w'(v), w'(v) + 1, … w'(v) + v − 1}
         Total miss trace = |W*|
         Remove all the nearest adjacent nodes, child
    nodes and the nearest adjacent nodes of the child
    nodes of t̄_{w'(v)}^{(4)}  from T̄
       end for
  end for
end while
```

To avoid repetitive trace selection, we implement a two-order sorting on all the candidate traces as Algorithm 1 demonstrates. For a node $\overline{t}_j^{(v)}$, its nearest adjacent nodes = $\left\{\overline{t}_{j-v}^{(v)}, \overline{t}_{j+v}^{(v)} \mid j - v \geq 0\ , j + v \leq W\right\}$. For example, as Figure 9 demonstrated, when a candidate trace $\overline{t}_{w'^{(v)}}^{(4)} = \overline{t}_4^{(4)}$ is observed in the tree $\overline{T}$, its nearest adjacent nodes $\{\overline{t}_0^{(4)}, \overline{t}_8^{(4)}\}$, child nodes $\{\overline{t}_4^{(2)}, \overline{t}_6^{(2)}, \overline{t}_4^{(1)}, \overline{t}_5^{(1)}, \overline{t}_6^{(1)}, \overline{t}_7^{(1)}\}$, and the nearest adjacent nodes of the child nodes $\{\overline{t}_2^{(2)}, \overline{t}_8^{(2)}, \overline{t}_3^{(1)}, \overline{t}_8^{(1)}\}$ are removed from $\overline{T}$.

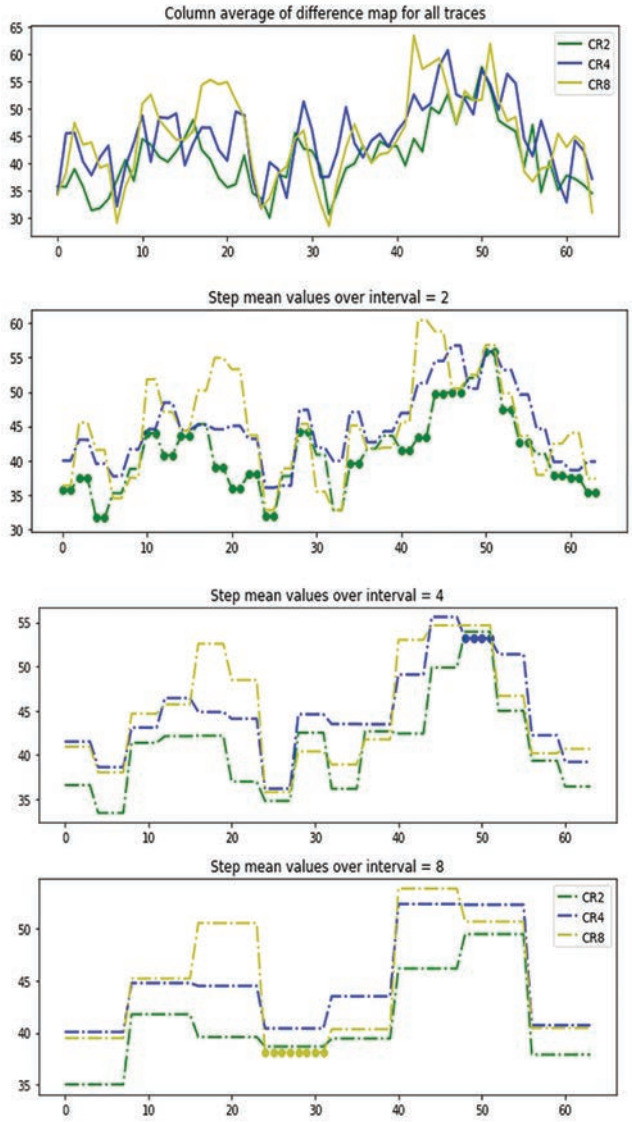

Figure 8. Example of column average of the difference map and its step mean over intervals. The *x*-axis stands for trace index, and the *y*-axis stands for the difference value over each trace. Dashed lines represent the step mean difference values over various observed intervals. Filled dots are the candidate traces.

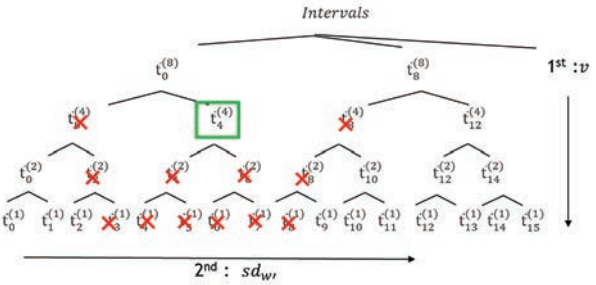

Figure 9. Binary search tree for all traces. For example, when a candidate trace (framed in green) is observed in the tree, its nearest adjacent nodes, child nodes, and the nearest adjacent nodes of the child nodes are removed from the tree (denoted by red x's).

The elements in the final set of W^* are the prospective missing traces, which are easily recovered by the GIN. We mention that our recommended sampling method does not consider the connected sampling crossing section width, and the effectiveness highly relies on the performance of the GIN.

Performance of recommended sampling. We have compared the reconstruction performance result of our recommended sampling survey with an average of 100 random samplings, and we report the improvement in Table 4.

Sampling recommendation on whole data set. We run the sampling recommendation on the same GoM data set in both inline and crossline directions. As the depth view in Figure 10 shows, the recommended sampling points are densely distributed in regions with lithologic features and sparsely distributed in channelized regions. This successfully captures the heterogeneity of the seismic image.

We note that our current recommendation of nonuniform sampling survey is given by the prior information of the

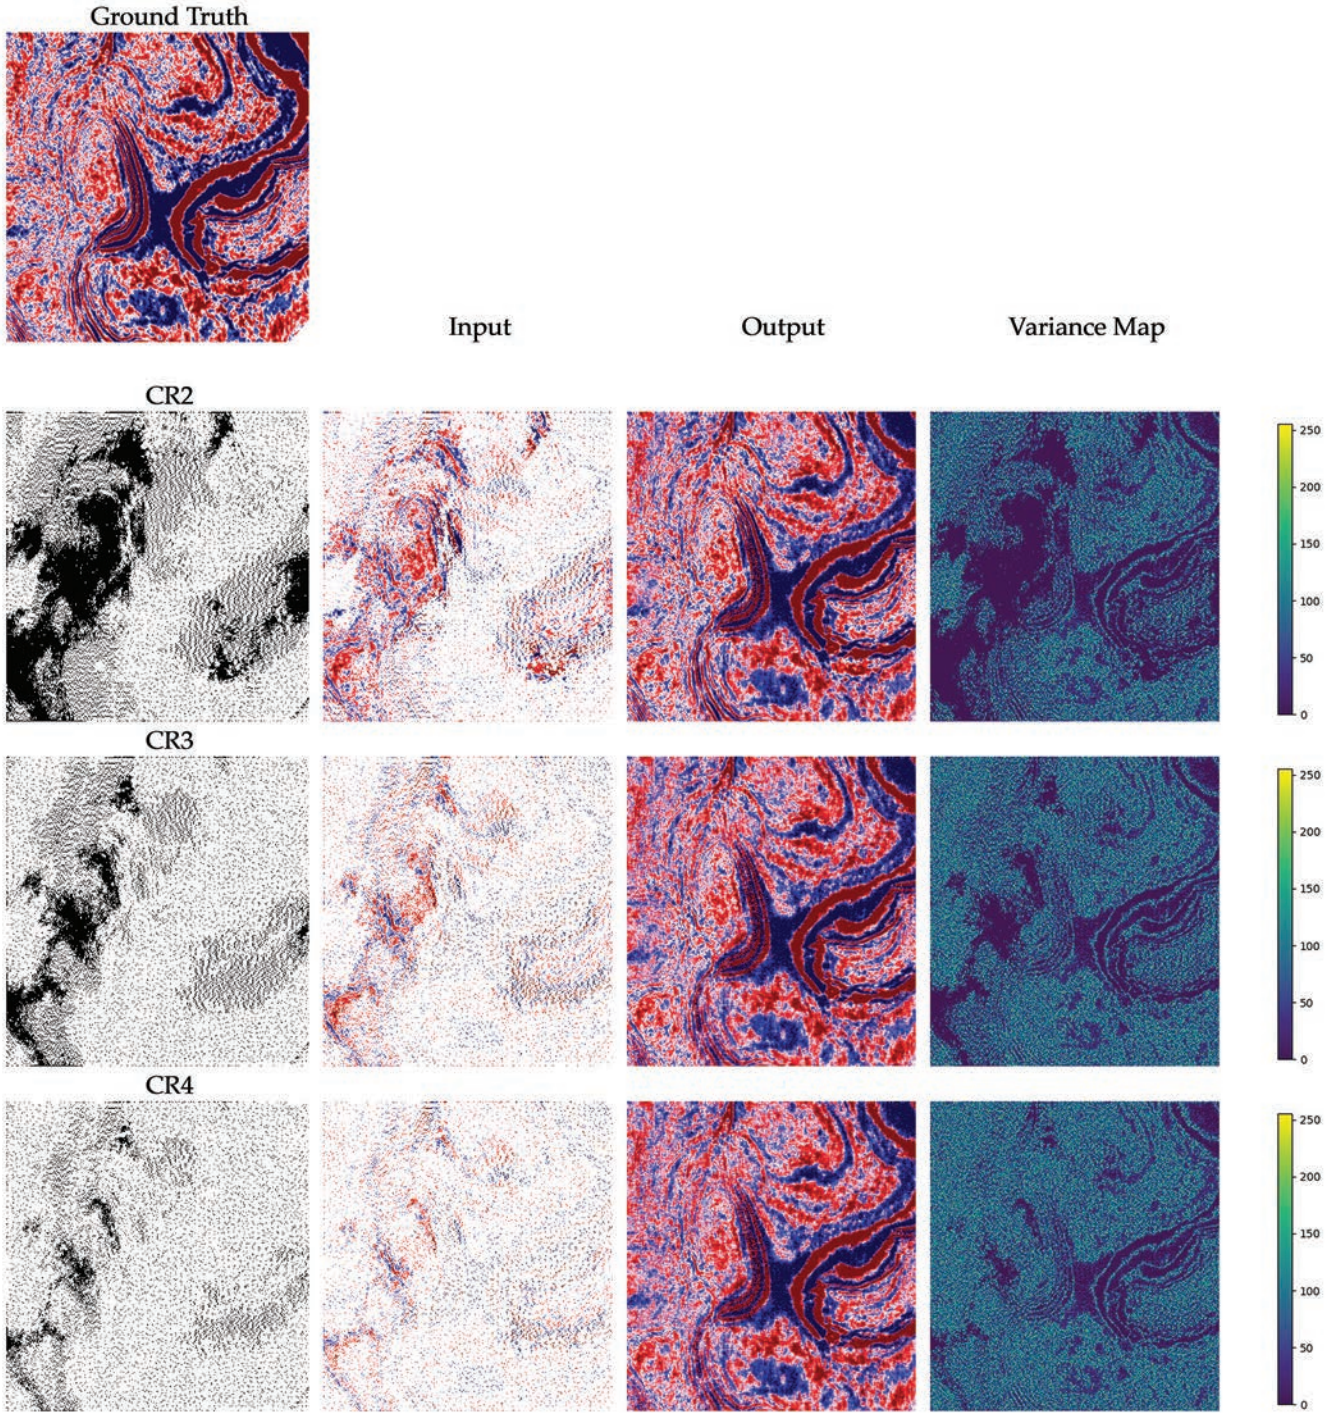

Figure 10. Stacked depth view comparing recommended sampling over CR = 2, 3, 4 on a 512 × 512 cropped region of the GoM data set. All the crossline and inline sampled images are stacked together to form a 3D reconstruction of the whole block, showing its 2D depth view at depth = 500 only for validation. (Seismic data courtesy of TGS)

historical images in the region. However, for new unexplored regions, the sampling could require additional geologic information.

Table 4. Performance comparison of recommended sampling and random sampling.

PSNR/SSIM	Recommended	Random
CR2	21.1 dB/0.79	15.2 dB/0.66
CR3	15.2 dB/0.59	12.2 dB/0.47
CR4	13 dB/0.46	10 dB/0.37

Conclusion

We designed and implemented a modification of the GIN model, the GIN-CS, and successfully tested its performance on uniform samplings with compression rates ×2, ×4, ×8, ×16. The GIN-CS demonstrates superior reconstruction performance relative to both the original GIN and the conventional biharmonic method. More precisely, we show that seismic imaging can be recovered successfully by filling the missing traces, revealing finer details even in cases of high compression rates. In addition, the proposed method runs approximately 300 times faster than the conventional biharmonic method. Finally, a strategy for constructing a recommendation of nonuniform survey is proposed for a field data set from the GoM, based on our results from a combination of limited amounts of uniform sampling experiments. TLE

Acknowledgments

The authors thank Anadarko for permission to publish this work and TGS for providing the seismic data set and approval for publication. We also appreciate the guidance of Cheng Zhan and Guillaume Richard from Anadarko in our research efforts.

Data and materials availability

Data associated with this content are confidential and cannot be released.

Corresponding author: xiaoyang.rebecca.li@gmail.com

References

Arjovsky, M., S. Chintala, and L. Bottou, 2017, Wasserstein GAN: arXiv preprint, arXiv:1701.07875.

Candès, E. J., J. K. Romberg, and T. Tao, 2006, Stable signal recovery from incomplete and inaccurate measurements: Communications on Pure and Applied Mathematics, **59**, no. 8, 1207–1223, https://doi.org/10.1002/cpa.20124.

Damelin, S. B. and N. S. Hoang, 2018, On surface completion and image inpainting by biharmonic functions: Numerical aspects: International Journal of Mathematics and Mathematical Sciences, **2018**, 1–8, https://doi.org/10.1155/2018/3950312.

Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, 2009, Imagenet: A large-scale hierarchical image database: IEEE Conference on Computer Vision and Pattern Recognition, https://doi.org/10.1109/CVPR.2009.5206848.

Donoho, D. L., 2006, Compressed sensing: IEEE Transactions on Information Theory, **52**, no. 4, 1289–1306, https://doi.org/10.1109/TIT.2006.871582.

Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, 2017, Improved training of wasserstein GANs. Proceedings of the 31st International Conference on Neural Information Processing Systems, 5769–5779.

Iizuka, S., E. Simo-Serra, and H. Ishikawa, 2017, Globally and locally consistent image completion: ACM Transactions on Graphics, **36**, no. 4, https://doi.org/10.1145/3072959.3073659.

Ledig, C., L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, et al., 2017, Photo-realistic single image super-resolution using a generative adversarial network: IEEE Conference on Computer Vision and Pattern Recognition, 4681–4690, https://doi.org/10.1109/CVPR.2017.19.

Li, C., C. C. Mosher, and S. T. Kaplan, 2012, Interpolated compressive sensing for seismic data reconstruction: 82nd Annual International Meeting, SEG, Expanded Abstracts, https://doi.org/10.1190/segam2012-1335.1.

Lu, P., M. Morris, S. Brazell, C. Comiskey, and Y. Xiao, 2018, Using generative adversarial networks to improve deep-learning fault interpretation networks: The Leading Edge, **37**, no. 8, 578–583, https://doi.org/10.1190/tle37080578.1.

Mandelli, S., V. Lipari, P. Bestagini, and S. Tubaro, 2019, Interpolation and denoising of seismic data using convolutional neural networks: arXiv preprint, arXiv:1901.07927.

Pathak, D., P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, 2016, Context encoders: Feature learning by inpainting: IEEE Conference on Computer Vision and Pattern Recognition, https://doi.org/10.1109/CVPR.2016.278.

Picetti, F., V. Lipari, P. Bestagini, and S. Tubaro, 2018, A generative-adversarial network for seismic-imaging applications: 88th Annual International Meeting, SEG, Expanded Abstracts, 2231–2235, https://doi.org/10.1190/segam2018-2995439.1.

Wang, X., K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y Qiao, and C. Change Loy, 2018, ESRGAN: Enhanced super-resolution generative adversarial networks: European Conference on Computer Vision, 63–79, https://doi.org/10.1007/978-3-030-11021-5_5.

Wang, Z., A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, 2004, Image quality assessment: From error visibility to structural similarity: IEEE Transactions on Image Processing, **13**, no. 4, 600–612, https://doi.org/10.1109/TIP.2003.819861.

Yu, J., Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, 2018, Generative image inpainting with contextual attention: IEEE Conference on Computer Vision and Pattern Recognition, 5505–5514, https://doi.org/10.1109/CVPR.2018.00577.