Geophysical Journal International

Geophys. J. Int. (2020) **223,** 1888–1898
Advance Access publication 2020 September 8
GJI Marine Geosciences and Applied Geophysics

doi: 10.1093/gji/ggaa422

Inpainting of local wavefront attributes using artificial intelligence for enhancement of massive 3-D pre-stack seismic data

Kirill Gadylshin,[1,2] Ilya Silvestrov[3] and Andrey Bakulin[3]

[1]*Institute of Petroleum Geology and Geophysics, pr. Koptyug* 3, 630090*, Novosibirsk, Russia. E-mail: gadylshin@gmail.com*
[2]*Novosibirsk State University, Pirogova* 2 *St.,* 630090*, Novosibirsk, Russia*
[3]*EXPEC Advanced Research Center, Saudi Aramco, Dhahran, Saudi Arabia*

Accepted 2020 September 1. Received 2020 August 17; in original form 2020 January 30

SUMMARY

We propose an advanced version of non-linear beamforming assisted by artificial intelligence (NLBF-AI) that includes additional steps of encoding and interpolating of wavefront attributes using inpainting with deep neural network (DNN). Inpainting can efficiently and accurately fill the holes in waveform attributes caused by acquisition geometry gaps and data quality issues. Inpainting with DNN delivers excellent quality of interpolation with the negligible computational effort and performs particularly well for a challenging case of irregular holes where other interpolation methods struggle. Since conventional brute-force attribute estimation is very costly, we can further intentionally create additional holes or masks to restrict expensive conventional estimation to a smaller subvolume and obtain missing attributes with cost-effective inpainting. Using a marine seismic data set with ocean bottom nodes, we show that inpainting can reliably recover wavefront attributes even with masked areas reaching 50–75 per cent. We validate the quality of the results by comparing attributes and enhanced data from NLBF-AI and conventional NLBF using full-density data without decimation.

Key words: Image processing; Neural networks; Numerical approximations and analysis; Seismic noise.

INTRODUCTION

In the early days of artificial intelligence (AI), it tackled and solved problems that are intellectually difficult for human beings, but relatively straightforward for computers. The real challenge to AI proved to be addressing the tasks that are straightforward for people to perform, but difficult for people to formally describe—problems that we solve intuitively (Goodfellow *et al.* 2016). Recently, deep learning approaches succeeded in image processing tasks such as object recognition, denoising, super-resolution (Halpert 2018), image inpainting (Liu *et al.* 2018), seismic trace interpolation (Mandelli *et al.* 2018; Wang *et al.* 2020) and seismic compressive sensing (Li *et al.* 2019).

In seismic imaging, the desire to improve reservoir description combined with advances in acquisition technology is pushing the amount of collected seismic data into the 'big data' category (Araya-Polo *et al.* 2017). Pre-stack data from modern high-density seismic surveys can reach hundreds and thousands of terabytes in size. Such data can often have a reduced signal-to-noise ratio, particularly in the case of single-sensor land seismic data, therefore, demanding new algorithms and approaches for its efficient processing and interpretation (Bakulin *et al.* 2018). One of the ways to leverage massive seismic data sets is to use pre-stack seismic attributes, namely some derivatives of the recorded data allowing us to represent it in a more compressed way as well as perform necessary manipulations, for instance, enhancement or interpolation. A classic example of such an attribute is a normal moveout velocity, which is usually picked or evaluated on a regular coarse grid, and then interpolated between the gridpoints serving as a critical ingredient to perform seismic imaging. Many novel approaches based on pre-stack data require estimation of multiple attributes for each particular point in the data cube, and their simultaneous calculation and usage is of great importance. Examples of such approaches are multidimensional data-driven local stacking and data enhancement techniques such as partial common-reflection surfaces (CRS) stack (Baykulov & Gajewski 2009), non-hyperbolic multifocusing (Berkovitch *et al.* 2011), and non-linear beamforming (NLBF, Bakulin *et al.* 2018, 2020). They all require the estimation of at least five wavefront attributes on a dense spatial and temporal grid, namely the 3-D *X*-*Y*-*T* grid. Direct estimation of wavefront attributes is usually the most time-consuming part of the entire workflow. In practical applications to real data, several challenges may require efficient interpolation of wavefront attributes. Most acute of these challenges are:

Acquisition geometry issues. 3-D seismic data may suffer from a significant irregularity of sources or receivers associated with existing infrastructure or environmental constraints. This results in:

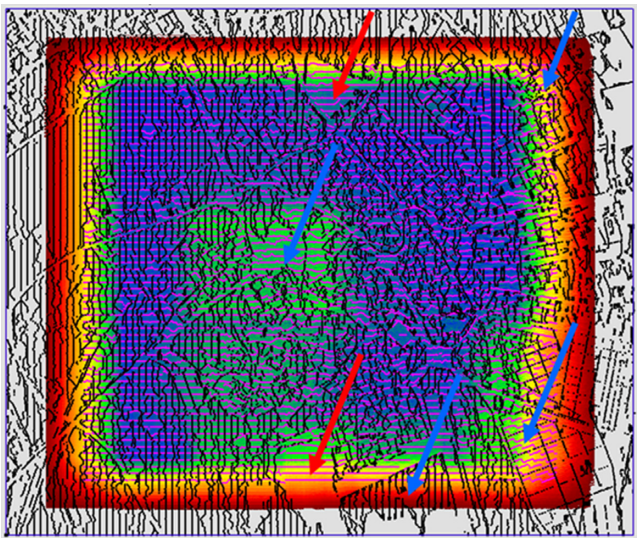

Figure 1. Acquisition geometry from a 3-D land seismic survey conducted in an urban area with restricted access: black dots denote vibrator positions, whereas purple ones show receiver locations. Observe overall irregularity (blue arrows) and significant large gaps (red arrows) that would lead to holes where wavefront attributes cannot be estimated reliably and would require interpolation.

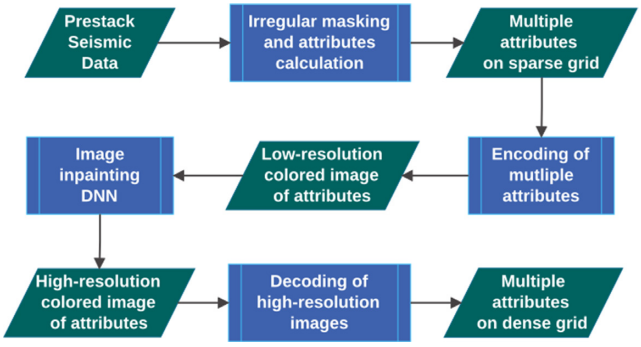

Figure 2. Flowchart of a proposed workflow. The parallelograms show input and output data, whereas the rectangles show different parts of the code/algorithm.

(i) Irregular/sparse distribution (Fig. 1, blue arrows), leading to unreliable estimation of wavefront attributes. Quality control approaches often flag spatial subvolumes where reduced density may render wavefront attributes below acceptable quality. Nevertheless, pre-stack data still require enhancement, and therefore interpolation of kinematic parameters from other reliable subvolumes may be the only option.

(ii) Significant gaps without any sources or receivers (Fig. 1, red arrows), leading to a similar effect along the edges. Enhancing data in the vicinity of such holes benefits from interpolation/extrapolation of wavefront attributes. Besides this, the reconstruction of missing data in such areas demands reliable wavefront attributes that can only be obtained by interpolation.

Data quality issues. Even when data is present, there are often spatial or temporal zones with high noise or 'bad' data that result in the unreliable estimation of wavefront attributes. Such flagged areas still contain seismic signals that could be salvaged using multidimensional stacking, provided a reliable estimate of the wavefront attributes is obtained. Interpolation of wavefront attributes is often a last resort. However, it should be able to handle the completely unpredictable spatial/temporal nature of such zones in the 3-D *X*-*Y*-*T* data grid.

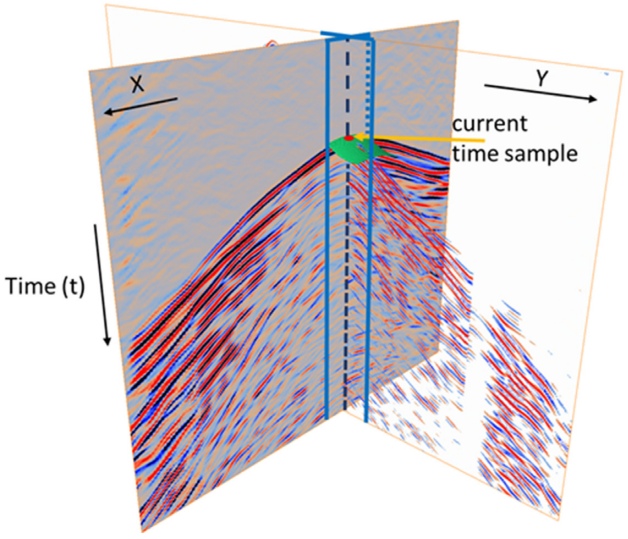

Figure 3. Wavefront attributes are estimated on a dense spatial-temporal grid within the entire volume of pre-stack seismic data. Local second-order traveltime surface (green surface) represents the estimated best fit moveout to actual seismic events observed at each temporal and spatial point of the grid (red dot). The surface is described by five coefficients representing wavefront attributes (dips and curvatures). The dashed dark blue line denotes a parameter trace in (*X*, *Y*, *T*) space where estimation is conducted. Parameter traces on a regular grid with defined spacing are required for data enhancement. Solid blue lines specify the extent of the rectangular estimation aperture.

Drive to computational efficiency. Even in the absence of two issues above, there is a constant need to improve computational efficiency for ever-increasing data volumes with high-density seismic data. One way to efficiently address it is by limiting computations of wavefront attributes to a reduced number of gridpoints followed by an interpolation to a complete volume. In contrast to uncontrollable effects of acquisition and data quality issues, excluded zones for performance improvement can be designed regularly or irregularly. Uniform exclusions or masking can benefit from established, simple methods such as bilinear interpolation.

In practice, all three factors above are simultaneously at play, with the first two factors being outside of our control. Quality control algorithms can flag problematic areas caused by acquisition and data quality issues ahead of the estimation process resulting in irregularly placed space–time masks scattered around the data volume. The resulting final mask is a superposition of masks from all three of the causes described. When the portion of the mask caused by the first and second factors is significant, then there is a need to handle the interpolation of irregularly masked data even if the regular mask was selected for improving compute efficiency in the third step. Under these circumstances, we focus here on designing and validating an efficient inpainting method that can handle the general case of interpolation in the presence of arbitrary non-uniform masking. To further stress-test the proposed algorithm, we deliberately focus on the escalating proportion of random irregular masks as a worst-case scenario. While regular-only masking for performance improvement may be less demanding, we believe that the new algorithm may have the potential to compete in computational

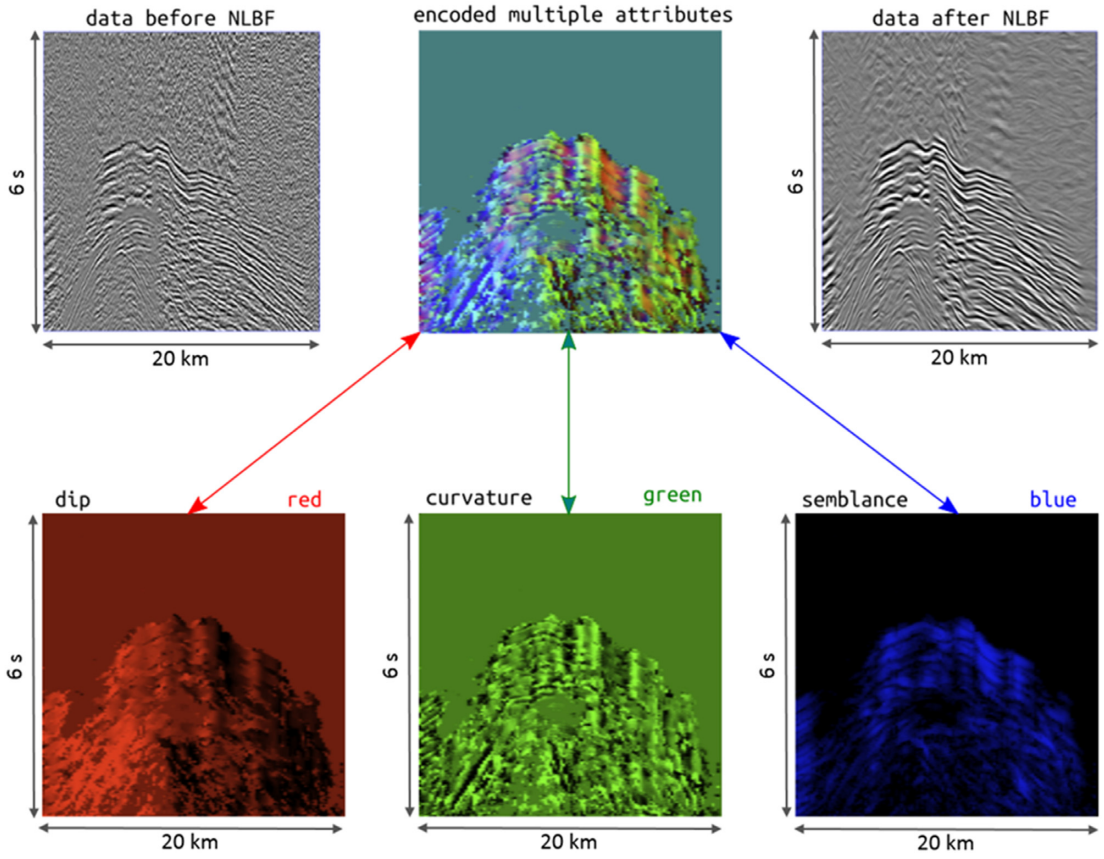

Figure 4. Example of colouring (encoding stage) for multiple kinematic wavefront attributes required for NLBF. Three quantities are encoded into a single RGB image comprised of three channels: red—dips, green—curvatures, and blue—semblance.

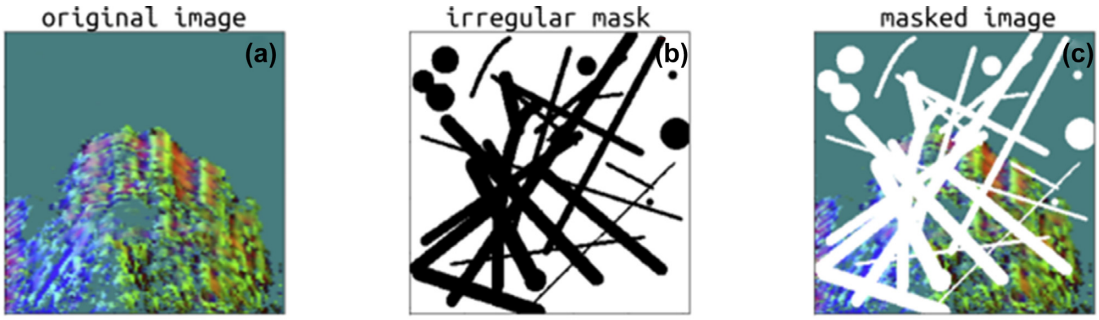

Figure 5. Original encoded RGB image of multiple attributes (a), randomly generated image mask (b) and an image with the mask applied (c).

efficiency with more straightforward interpolation approaches applicable in this specific case. However, this point is left beyond the scope of this study.

OVERVIEW OF APPLICABLE INTERPOLATION METHODS

Inpainting is the process of reconstructing lost or deteriorated parts of images and videos. It is also known as image interpolation and refers to the application of sophisticated algorithms to replace corrupted parts of the image data. All the existing techniques use information from the available known image areas to fill the gap. In the commonly used OpenCV (open computer vision) library, two different algorithms are popular. The first of these algorithms is based on a fast marching method (Telea 2004). The algorithm starts from the boundary and goes inside the region of interest, gradually filling all unknown regions. Another approach is based on fluid dynamics and utilizes partial differential equations (Bertalmio *et al.* 2001). The basic principle is heuristic: travelling first along the edges from known regions to unknown regions. The algorithm continues isophotes while matching gradient vectors at the boundary of the inpainting region. Missing colour is filled to reduce the minimum variance in that area.

Along with the classical interpolation techniques, DNN-based image super-resolution approaches demonstrate superior performance and outcomes (Ledig *et al.* 2017). Recent deep learning approaches led to significant advances in image inpainting. A majority of these methods focus on inpainting inside rectangular regions located around the centre of the image and often rely on expensive post-processing (see Liu *et al.* 2018, and references therein). DNNs

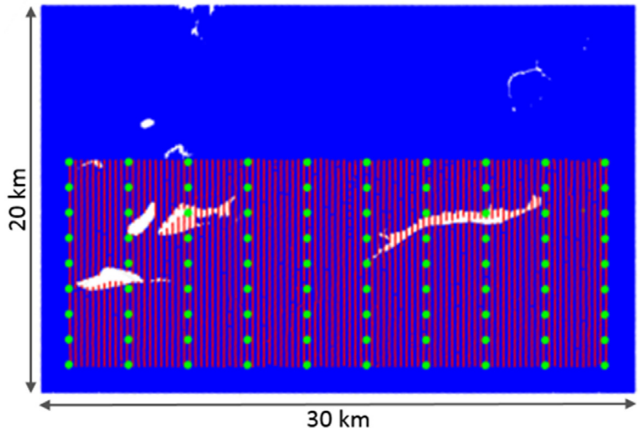

Figure 6. Data acquisition geometry for the marine data sets recorded with ocean bottom nodes. Shot positions are shown in blue, and receivers are in red. Green points designate selected locations identifying common-receiver gathers that were included in a limited training data set (0.8 per cent of the total data size).

use convolutional filters on images, replacing the removed content with a fixed value, which is often dependent on the initial values used. To properly handle irregular masks, Liu *et al.* (2018) proposed using partial convolutions, where the convolution is masked and renormalized to be conditioned only on valid pixels. This model outperforms other methods for irregular masks. In this study, we evaluate applications of the DNN based on partial convolutions to the problem of inpainting of local wavefront attributes.

DNN-BASED INPAINTING OF WAVEFRONT ATTRIBUTES FOR NLBF

General workflow

To address the interpolation challenge in the context of data enhancement, we pursue a novel approach referred to as NLBF-AI. First, we evaluate the wavefront attributes inside unmasked areas using the conventional estimation approach. Then, we reconstruct or inpaint remaining attributes using a leading-edge machine learning (ML) technique—a DNN with partial convolutional layers (Liu *et al.* 2018). This inpainting is based on training a DNN to provide high-resolution output for a given low-resolution input. We use 'low-resolution' as a general term designating a sparse volume with masked areas irrespective of the actual spatial size and distribution of masked elements. Similarly, 'high-resolution' simply denotes a densely filled volume with attributes available at all points. Inpainting is significantly more computationally efficient than the standard calculation of attributes at each location, whereas it is fully capable of capturing a sufficient level of detail required for data enhancement. A schematic workflow is presented in Fig. 2.

First, we assemble a cumulative mask caused by the effects of acquisition geometry, data quality issues, as well as exclusion areas deliberately created to increase computational performance. We then proceed with computing wavefront attributes in unmasked areas from pre-stack seismic data. Grids with random exclusion areas are well suitable for efficient sampling and capturing of the main wavefront features present for specific geologies within the volume of interest. We associate these computed multiple attributes with coloured images using a particular encoding scheme transforming multiparameter attributes into a single point of a coloured image. Once we obtain the encoded images containing desired attributes, the trained DNN based on partial convolutions inpaints the

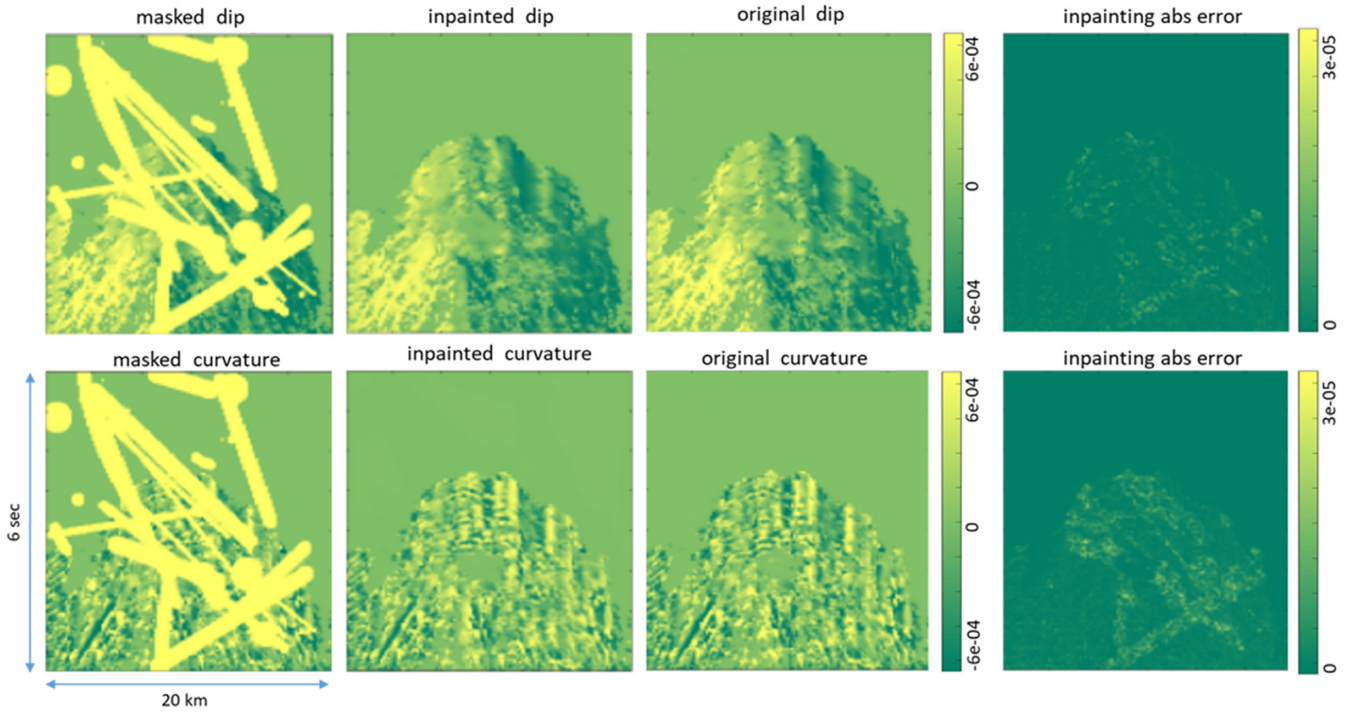

Figure 7. Evaluating multiparameter wavefront attributes predicted by trained DNN in a single vertical plane. From left- to right-hand panel: masked, predicted and original attributes, along with the difference between predicted and original attributes plotted using the same scale (last column). The top row represents dips, whereas the bottom row shows curvatures.

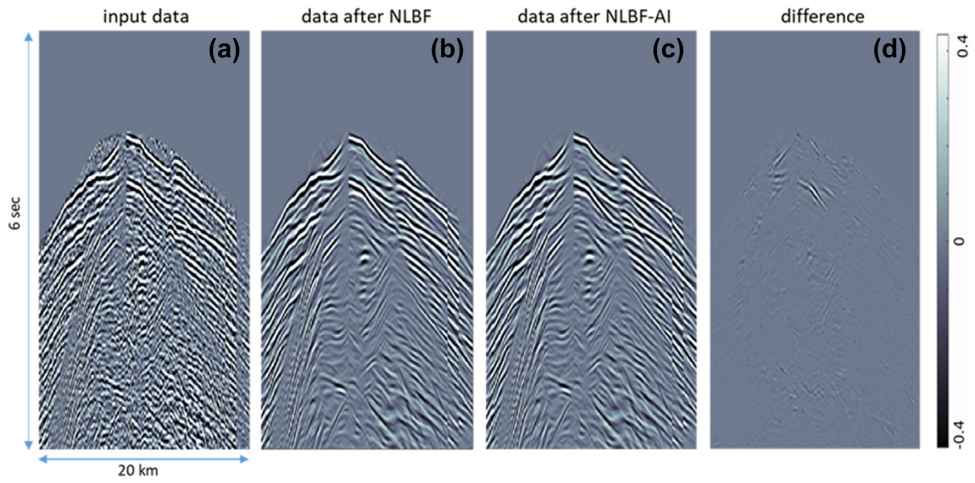

Figure 8. Evaluating enhanced data: input data (a), conventional NLBF (b), NLBF-AI (c) and the difference between the two versions of NLBF data (d).

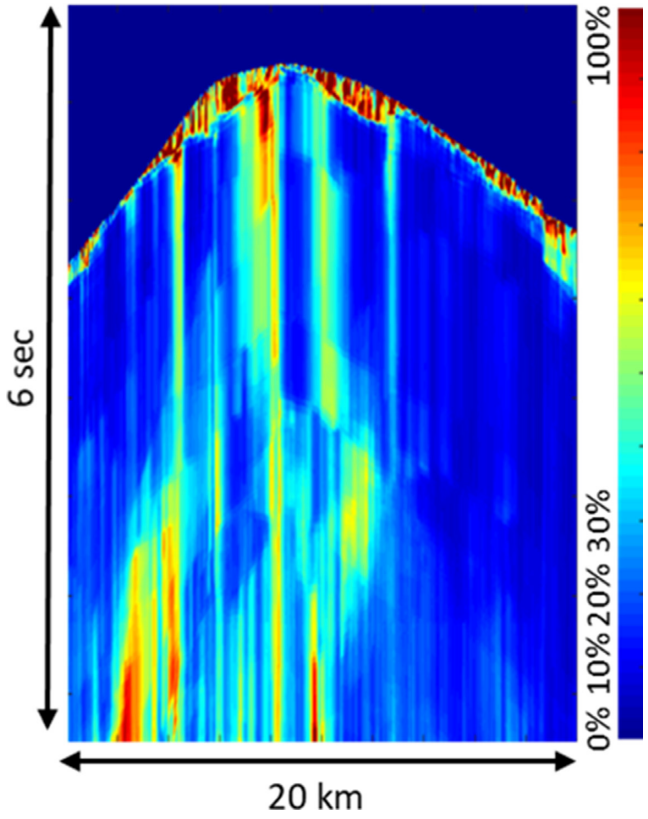

Figure 9. NRMS between two versions of the gather enhanced using wavefront attributes computed on the full dense grid (standard NLBF) and interpolated with the inpainting (NLBF-AI). NRMS is calculated trace by trace using a sliding window of 150 ms. The average NRMS for the entire ensemble is around 21 per cent indicating that gathers are nearly identical for processing purposes.

remaining masked regions and predicts a complete high-resolution coloured image (filling gaps or holes). The last step is decoding the predicted coloured image back into multiple attribute spaces resulting in a dense inpainted grid over the entire volume as required for enhancement.

Non-linear beamforming

As an example, we illustrate the application of this approach for the estimation of local wavefront attributes used in non-linear time-delay beamforming (Bakulin *et al.* 2020). NLBF belongs to a family of multidimensional stacking approaches (references) that are critical for pre-stack enhancement of data with the low signal-to-noise ratio, as demonstrated on many practical examples by Bakulin *et al*. (2018, 2020). However, challenges above associated with wavefront attributes estimation often limit NLBF applicability. We aim to address them through a novel inpainting approach. To understand this novel implementation, it is essential to remind the algorithmic details of NLBF-based data enhancement. This method comprises a local summation of nearby traces after application of time shifts and formally can be written as:

$$u\ (x_0, y_0, t_0) = \sum_{x \in B_0} w\,(x, y)\, u\,(x, y, t_0 + \Delta t)\,, \tag{1}$$

where $u(x, y; t)$ represents a trace with coordinates x and y. The coordinates of the output trace after beamforming are given by x_0, y_0. The summation is performed over a local region B_0 (defined by summation aperture) around the output trace along a traveltime surface with some moveout Δt. We assume that a second-order surface can locally approximate wavefront with the following moveout:

$$\Delta t = t\,(x, y) - t_0\,(x_0, y_0) = A\Delta x + B\Delta y + C\Delta x\Delta y + D\Delta x^2 + E\Delta y^2, \tag{2}$$

where A, B, C, D and E are unknown wavefront attributes that also serve as beamforming coefficients; Δx and Δy represent spatial shifts of the summed trace to the output trace.

The first step in the NLBF data enhancement procedure is the estimation of local kinematic attributes that describe coherent local events. These parameters are usually estimated using an operator-oriented approach (Hoecht *et al.* 2009). In this scheme, distinctive auxiliary parameter traces are introduced, and the estimation of wavefront attributes is done in samples of parameter traces using all actual data traces within an area defined by the estimation aperture (Fig. 3).

The unknown coefficients A, B (first-order derivatives of the wavefront, or dips) and C, D, E (second-order derivatives, or curvatures) are estimated by scanning many different beamforming surfaces and finding one with the best coherency characterized by the maximum value of a semblance function, S. Optimization can

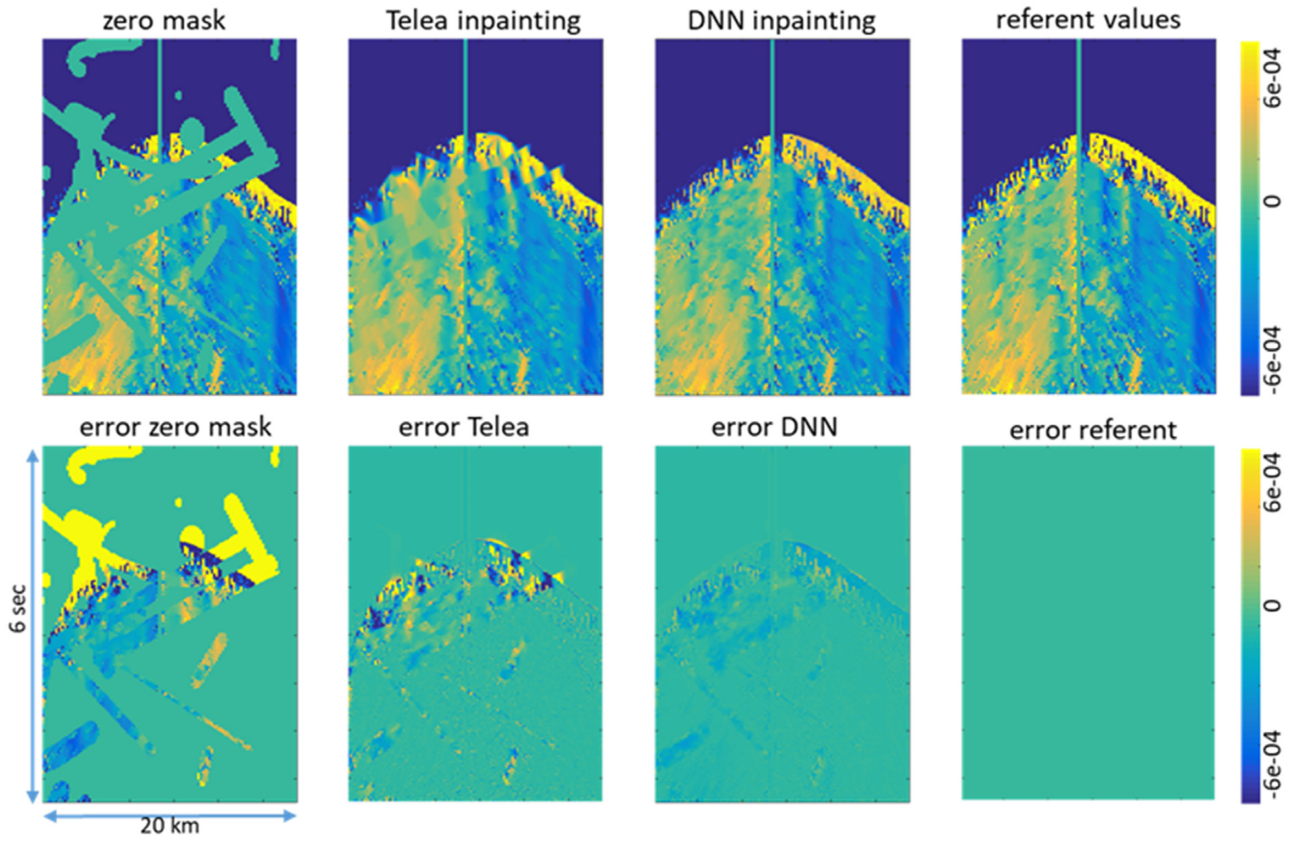

Figure 10. Various versions of wavefront attribute A (dip) reconstructed by different methods (from left- to right-hand panel): zero-value infill, interpolation using Telea's algorithm, DNN inpainting, brute-force calculation using '2 + 2 + 1' scheme from standard NLBF (reference). The second row shows images with the corresponding errors calculated as a straight difference between the reconstructed and reference values.

be implemented either by simultaneously searching for all five parameters, or sequentially searching for one parameter after another as in coordinate-descent method. Current practical implementations of NLBF (Bakulin *et al.* 2020) are based on hybrid '2 + 2 + 1' strategy, searching first for parameters A and D in the $X - t$ plane, than B and E in the $Y - t$ plane, and finally finding parameter C in the X- Y plane with the first four parameters being fixed.

AI-assisted NLBF (NLBF-AI)

Following in the footsteps of the '2 + 2 + 1' strategy, let us first outline the estimation step restricted to $X - t$ plane only. Only the dip A, curvature D, and semblance S are evaluated using a subset of the data and output at each time sample of parameter traces:

$$\Delta t = t\,(x) - t_0\,(x_0) = A\Delta x + D\Delta x^2. \tag{3}$$

For a fixed point (t, x) of the regular grid, our encoding stage transforms a triplet of multiparameter attributes $\langle A, D, S\rangle$ into a coloured image pixel $\langle Red, Green, Blue\rangle$ (RGB). As a result, local wavefront attributes A, D, and semblance S in each 2-D plane are represented now as a single coloured image (Fig. 4). We store RGB images without compression, that is we are still using floating-point values in the three channels. By design, we have a bijection (one-to-one correspondence) between coloured image and pair of local kinematic attributes along with the associated semblance.

We further assume that the three factors described above resulted in a masked area where wavefront attributes have to be interpolated. Intuitively, it is clear that interpolation may become harder with the increasing volume of the masks. We quantitatively evaluate the quality of interpolated attributes and enhanced data as a function of the escalating mask proportion. To stress-test the algorithm and maintain generality, we generate irregular masks as a random combination of geometrical objects like circles, straight lines, and sectors of an ellipse with random coordinates, sizes and inclinations (Fig. 5). The only parameter that is predefined by the user is a mask area as a percentage of the total image. In all points (t, x) where the mask is applied, wavefront attributes are interpolated, and computation is avoided. If only a third factor of computing efficiency is of the essence, then the area of the mask as a percentage of the total area represents the total computational savings. Then, we construct the convolutional DNN with a U-Net like architecture (Ronneberger *et al.* 2015) using partial convolutions (Liu *et al.* 2018). To properly handle irregular masks, the usage of partial convolutional layers is essential. The technical details related to the suggested DNN are presented in the Appendix.

Full enhancement of pre-stack 3-D data requires the estimation of dips and curvatures in two orthogonal planes ($X - t$ and $Y - t$). The hybrid '2 + 2 + 1' approach estimates A and D in one plane, B and E in another, and cross-coupling parameter C during the final step. To extend the above method from a single 2-D plane, we note that seismic offset gathers behave similarly in both orthogonal planes. Indeed, each plane contains a 2-D snapshot or cross-section of the same original 3-D wavefield recorded along a specific azimuth

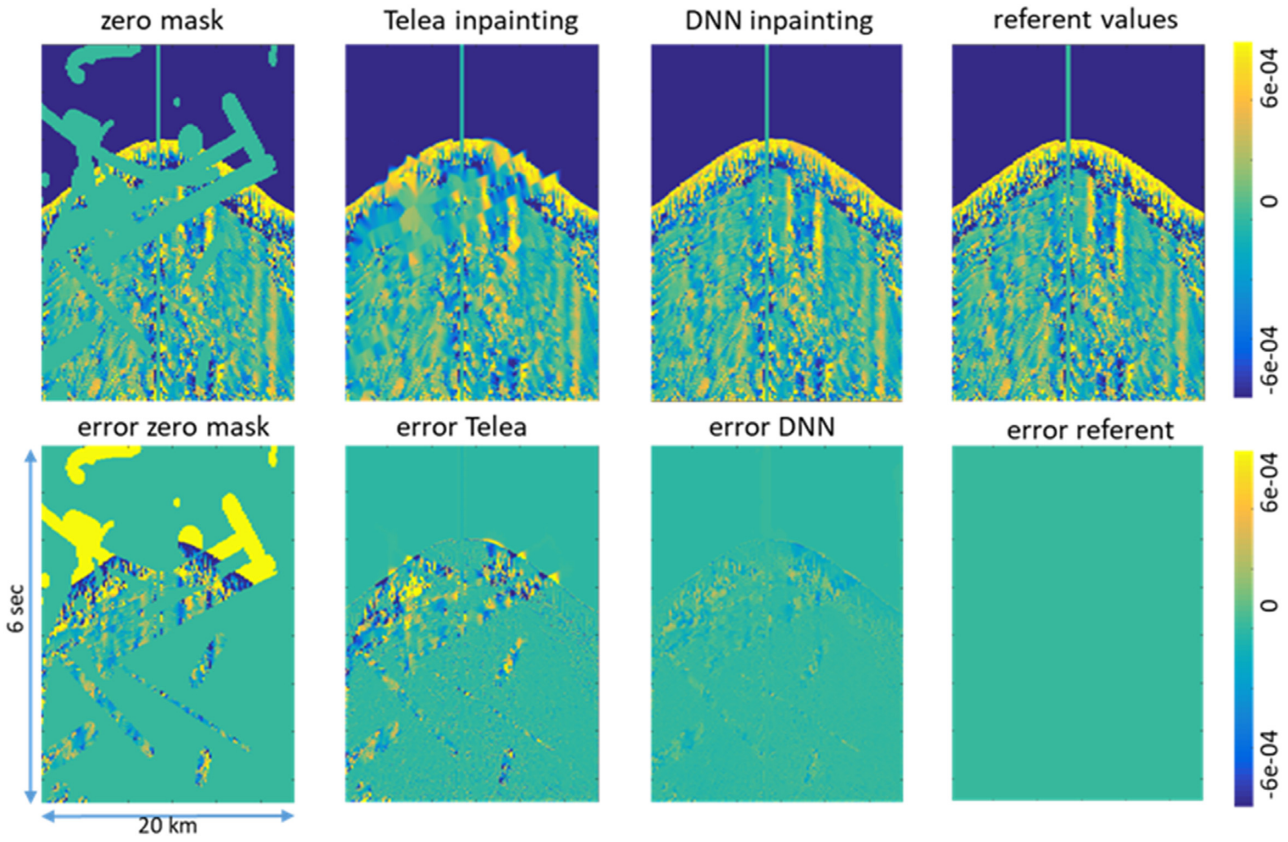

Figure 11. Same as Fig. 9, but for parameter D.

inside the subsurface. As a consequence, the inpainting workflow can handle orthogonal offset planes simultaneously. To achieve this, we perform a single DNN training using representative gathers mixed from both planes and then apply identical predictions to inpaint dips and curvatures in each cross-section. As a result, we obtain estimates of A, B, D and E, whereas the remaining parameter C is separately evaluated at the end using conventional means.

APPLICATION TO A MARINE OBN DATA SET

DNN training

The crucial step in deep learning is the training of the neural network. The first question to answer whether we should use pre-trained DNN for all types of seismic data, or should we retrain our network for every particular type of the data or even for every specific geographic region? The answer depends on many factors, such as training time as compared to the local kinematic attributes estimation time. Another consideration is a dependence on specific features of the seismic data that may be present in some areas and absent in the others (overfitting related to the non-representativeness of the training data set).

In our field case study, we use a marine data set recorded with OBNs and airguns with a pre-stack volume of around 15 TB. The data consists of 11 027 common-receiver gathers. The OBN survey size is 20 km $\times$ 31 km, comprising of 95 receiver lines. Enhancement of data with a low signal-to-noise ratio in the entire offset range from zero to 20 km is a geophysical must for velocity model building, specifically to achieve reliable first-break picking for tomography and robust waveforms for FWI.

To make out method practical on giant seismic volumes, the DNN is trained only on a small subset of data shown in Fig. 6. We utilize 1 per cent of all common-receiver gathers and then calculate the kinematic wavefront attributes using conventional NLBF. Calculated attributes are converted into a collection of 'RGB' images. These data are randomly split into a training data set (80 per cent of images), and a validation data set (20 per cent). To achieve diversity and representativeness, common-receiver gathers are extracted from an evenly spaced nearly regular subgrid (Fig. 6).

The training is performed on a modern high-performance computing cluster Ibex from KAUST (https://www.hpc.kaust.edu.sa/ibex) using a single-node with four Tesla GPUs P100. Compute time used for training is negligible compared to the calculation time required for the estimation of kinematic attributes for the whole data set. To avoid overfitting, we use validation-based early stopping regularization techniques. The error on the validation data set is used as a proxy for the generalization error in determining when overfitting has begun.

The trained network inpaints encoded wavefront attributes delivering complete high-resolution images without holes. Applying a decoder to these images finalizes the workflow and provides required wavefront attributes (A, B, C, D and E) at every point of the pre-stack estimation grid in space and time.

Verification through comparison of standard and AI-assisted NLBF

To evaluate the quality of inpainted attributes, we create a testing data set consisting of randomly positioned common-receiver gathers that are different from the training and the validation data set. The neural network has not seen wavefront attributes estimated from

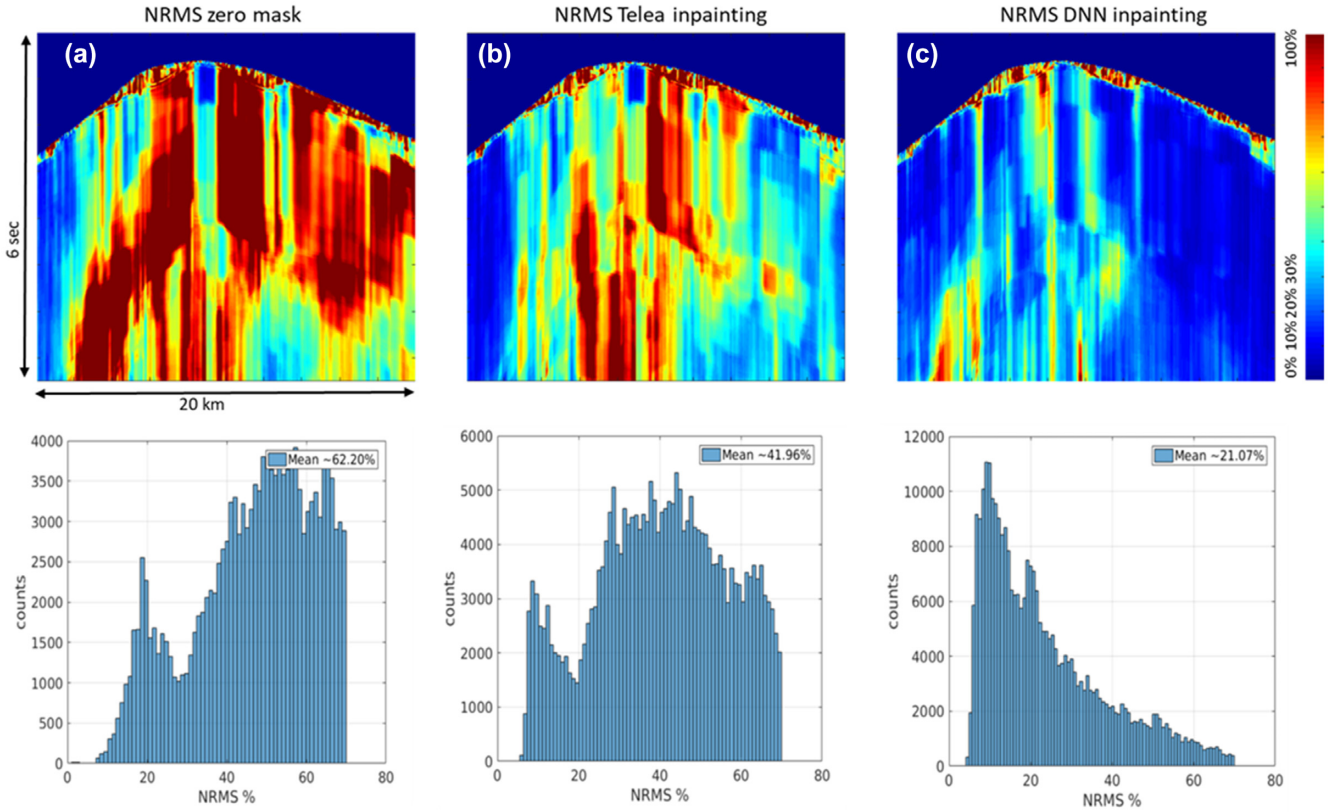

Figure 12. Similarity/repeatability quantified as NRMS between data enhanced with standard NLBF and NLBF using various methods of attribute interpolation. The top row displays NRMS computed in a sliding window, whereas the bottom row shows a histogram showing NRMS distribution for: (a) zero-value infill; (b) Telea inpainting and (c) DNN inpainting. DNN leads to the smallest NRMS values demonstrating the closest similarity to the reference and indicating better interpolation quality.

these gathers before. We calculate the error between kinematic attributes obtained with the conventional method and DNN inpainting using the Frobenius norm. We find that the average accuracy of reconstructed kinematic attributes using a testing data set is ∼95 per cent. Fig. 7 confirms that infilled values using the neural network are similar to the attributes calculated using the traditional estimation scheme of the NLBF algorithm.

Since attribute estimation is an intermediate step of NLBF enhancement, it is desirable to quantitatively evaluate how the inpainting of attributes affects data after beamforming (Fig. 8). We use normalized RMS (NRMS) to measure repeatability or similarity between two versions of enhanced data with and without interpolation. NRMS is a rigorous sample-by-sample metric used for evaluating repeatability between two data sets in 4-D seismic (Kragh & Christie 2002). The NRMS varies in the range of [0,200] per cent where the value of 0 per cent indicates perfect repeatability (identical traces), and 200 per cent corresponds to uncorrelated traces. In the presence of strong reservoir signals, NRMS on the order of 20–40 per cent often remains acceptable for this exceptionally accurate task of identifying minute 4-D reservoir differences. If such differences are acceptable in 4-D seismic, we postulate that, in the context of seismic processing for exploration, similar values of NRMS would imply that two processing results must be nearly identical for a practical purpose. The NRMS plot (calculated trace by trace using a sliding window with a 160 ms length) between two versions of the data, enhanced using a fully computed dense grid of wavefront attributes and one infilled with the inpainting, is shown in Fig. 9. The average NRMS of 30 per cent suggests a very close similarity between the two data. Zooming into local details, in the areas with a high signal-to-noise ratio (coherent signal dominates), we achieve almost perfect repeatability or low NRMS (dark blue colour). In contrast, areas with a low signal-to-noise ratio (random noise dominates) exhibit low repeatability or high NRMS (red colours). Such behaviour is expected for areas dominated by noise since semblance has no clear global maxima due to a lack of coherent events. Instead, there are many small local maxima due to noise, thereby making difficult any stable estimation of the wavefront attributes.

Quality comparison with other interpolation approaches and the effect of the total percentage of the masked areas

To evaluate the quality of interpolation, we compare DNN inpainting with commonly used fast marching inpainting (Telea 2004) implemented in the Open Source Computer Vision Library (OpenCV, https://opencv.org). The reconstructed parameters are shown in Figs 10 and 11. We also compare the results with a most straightforward zero-value infill that simply assigns zero values to dips and curvatures inside the interpolated areas. DNN inpainting provides better results compared to fast marching, which, in turn, outperforms zero-value infill. As before, we are interested in data reconstruction quality after beamforming, so we calculate NRMS between the data enhanced using attributes from these three scenarios

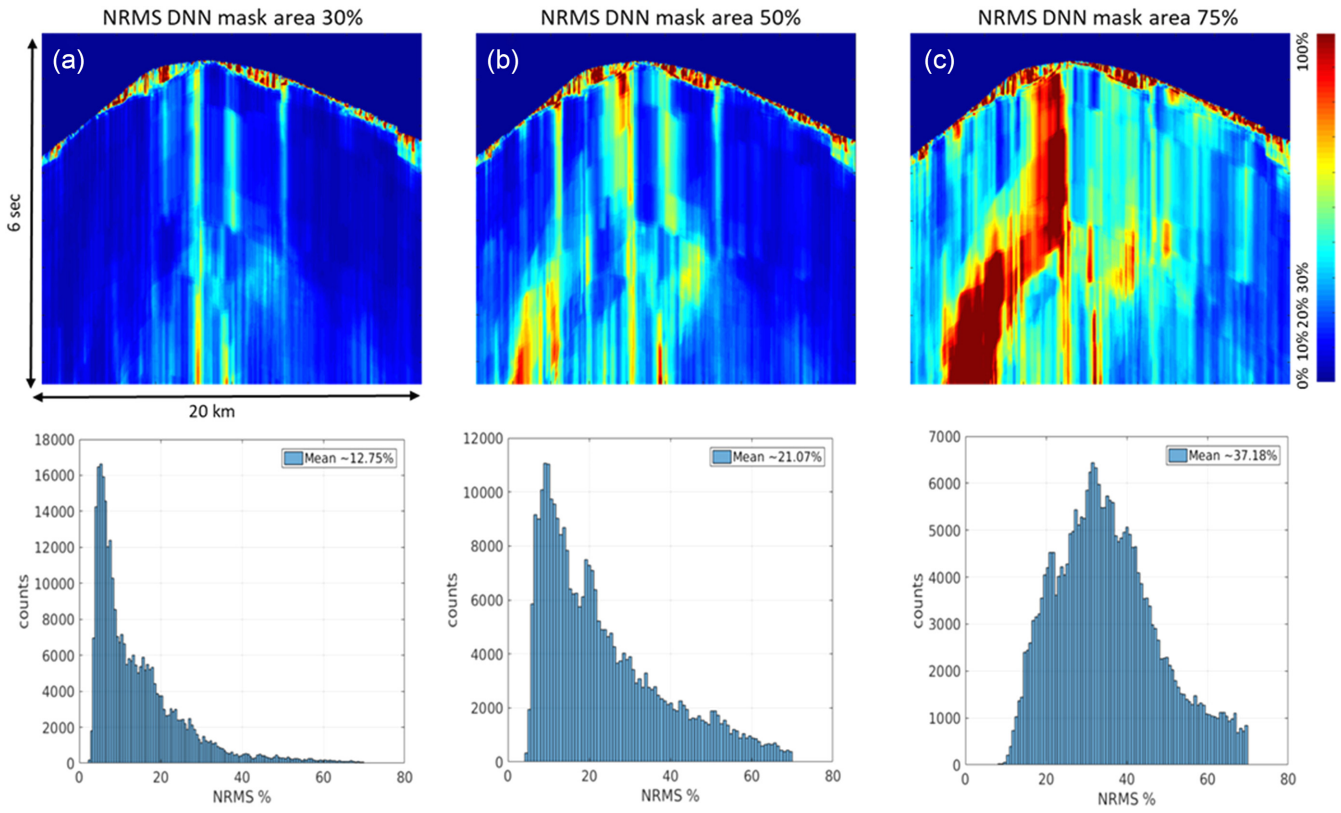

Figure 13. Effect of the total percentage of the masked area on enhanced data. Data enhanced with standard NLBF is used as a reference so that NRMS plots and histograms are computed for the difference between conventional NLBF results and NLBF-AI with the mask of 30 per cent (a), 50 per cent (b) and 75 per cent (c), respectively. A more significant percentage of masked areas leads to increased NRMS and reduced similarity. Nevertheless, mean values of 13, 21 and 37 per cent suggest that all data versions are nearly identical for seismic processing purpose to one obtained with standard NLBF.

and original data enhanced by NLBF (Fig. 12). The average NRMS between NLBF and NLBF-AI using inpainting is only 21 per cent. Even from a strict point of view of 4-D seismic, such repeatability indicates excellent agreement despite the total percentage of the masked area being relatively high at 50 per cent. In contrast, average NRMS values for zero-value infill and Telea inpainting are higher, ∼62 and ∼41 per cent, respectively. We conclude that the proposed AI approach delivers enhanced data that is measurably closer to the desired reference obtained without parameter interpolation.

How much decimation can be tolerated before there is a significant impact on data quality? To answer this question, we evaluate three different DNN inpainting scenarios with mask areas covering 30, 50 and 75 per cent of the total area, respectively. The corresponding NRMS plots and histograms are presented in Fig. 13. The average NRMS values are ∼13, 21 and ∼37 per cent, for the total percentage of masked areas 30, 50 and 75 per cent, respectively. While increasing NRMS indicates less and less similarity, the actual values are still acceptable for exploration processing. Going above 75 per cent leads to an escalating NRMS, indicating loss of resolution on wavefront attribute estimation. We conclude that masking of 50–75 per cent of the attributes on offset plane images represents the optimal range for computational performance: it delivers a significant speedup factor of 2–3*x* while maintaining reasonable similarity (NRMS ∼ 20–40 per cent) to original data, enhanced without interpolation.

DISCUSSION

We have shown that NLBF-AI with embedded inpainting effectively addresses practical need to interpolate wavefront attributes within large volumes with irregular gaps caused by acquisition geometry and data quality issues. In data enhancement with standard NLBF, the estimation phase consumes more than 90 per cent of the computational cost. In NLBF-AI, we have the flexibility to split the estimation phase between inpainting and conventional brute-force attribute calculation. Inpainting computing time is virtually negligible compared to traditional estimation time. As a result, performance improvement can be achieved that is proportional to the masked volumes. Creating additional controlled masking, we could restrict conventional estimation to an even smaller volume and let intelligent DNN inpainting to reconstruct parameters in the remaining parts. In the presented example with random masking of 50 per cent, we accomplish the speedup factor of two. Overall masked volume is always a combination of subvolumes masked due to acquisition geometry, data quality as well as exclusion areas aimed to reduce computation. At the same time, we have demonstrated that enhanced data with standard NLBF and NLBF-AI are nearly identical for processing purposes. While more standard interpolation methods can deliver quality similar to inpainting with the periodic placement of exclusion zones, they are prone to failure when masked areas become substantial or experience irregularity. At the same time, the computational cost of the inpainting approach is affordable and approaching the cost of these standard methods.

Since regularity of exclusion zones cannot be guaranteed on real field data (due to the unpredictable nature of acquisition geometry and data quality issues), we reason that using a more general inpainting approach in all cases could make good practical sense. In case of deliberate masking of 50–75 per cent of the volumes for computational purposes, the achieved speedups have a material impact on reducing processing turnaround of massive modern pre-stack seismic data sets with sizes comprising hundreds and thousands of terabytes. Periodic decimation could be another option that is easily handled by all methods inpainting included.

While an application of the inpainting to seismic data enhancement with NLBF-AI is compelling, similar attribute interpolation schemes can be applied to a vast collection of problems requiring the estimation of seismic attributes on massive volumes of pre-stack seismic data. Possible examples include the estimation of isotropic or anisotropic velocity fields, parameters for CRS stack, multifocusing parameters, and coefficients for isotropic and anisotropic amplitude vs offset analysis. Likewise, in the area of seismic monitoring, where multiple images are obtained in calendar time, and corresponding attributes are extracted many times over, there is a potential additional benefit of interpolating these parameters with deep learning without repeating expensive processing and even acquisition.

CONCLUSIONS

Local wavefront attributes form a skeleton of multidimensional pre-stack data and localize dominant coherent local events. Such a kinematic skeleton creates a foundation for powerful processing steps such as data-driven enhancement and interpolation. For example, powerful enhancement with non-linear beamforming requires an estimation of five wavefront attributes at every point on a dense 3-D *X*-*Y*-*T* grid. We present an advanced non-linear beamforming method (NLBF-AI) for data enhancement that combines conventional compute-intensive wavefront attribute estimation with a cost-effective DNN-based inpainting interpolation. Such a combination allows addressing undesirable irregular gaps in wavefront volumes caused by acquisition geometry and data quality issues as well as intensional exclusions for improving computational performance. In standard NLBF, estimation is conventionally performed using costly local or global optimization. Acquisition geometry gaps and data quality issues may lead to significant irregular exclusion zones in wavefront attribute volumes. Besides, the brute-force estimation phase consumes more than 90 per cent of the computing effort of overall data enhancement with standard NLBF. NLBF-AI represents an advanced version combining conventional estimation with an inpainting method based on DNN that can reliably and efficiently interpolate 5-D wavefront attributes inside irregular areas excluded from conventional estimation. Inpainting can handle large holes and unevenly spaced exclusion areas of arbitrary shapes and forms. We have applied NLBF-AI to the real OBN data set with a challenging signal-to-noise ratio and compared it with standard NLBF. We have shown that even when inpainting has to interpolate and infill intentionally excluded areas as large as 50–75 per cent, it still delivers attributes and enhanced data of almost the same quality as conventional NLBF as judged by strict NRMS metrics. For irregular exclusion zones, inpainting delivers better quality than other methods not based on AI. Considering the low computational cost of inpainting and high quality of interpolated attribute interpolation, we envision increased use of NLBF-AI in seismic processing of huge seismic volumes.

ACKNOWLEDGEMENTS

The authors would like to thank Maxim Dmitriev (Saudi Aramco) for assistance with the field study. One of the authors (Kirill Gadylshin) was partially supported by the grant from thePresident of the Russian Federation for young scientists – MK-670.2019.5. . The real data set associated with this study is confidential and cannot be released.

REFERENCES

Araya-Polo, M., Dahlke, T., Frogner, C., Zhang, C., Poggio, T. & Hohl, D., 2017, Automated fault detection without seismic processing, *Leading Edge,* **36**(3), 194–280.

Bakulin, A., Silvestrov, I., Dmitriev, M., Neklyudov, D., Protasov, M., Gadylshin, K. & Dolgov, V., 2020. Nonlinear beamforming for enhancement of 3D pre-stack land seismic data, *Geophysics,* **85,** V283–V296.

Bakulin, A., Silvestrov, I., Dmitriev, M., Neklyudov, D., Protasov, M., Gadylshin, K., Tcheverda, V. & Dolgov, V., 2018. Nonlinear beamforming for enhancing pre-stack data with challenging near surface or overburden. *First Break,* **36**(12), 121–126.

Baykulov, M. & Gajewski, D., 2009. Pre-stack seismic data enhancement with partial common-reflection-surface (CRS) stack, *Geophysics,* **74**(3), V49–V58.

Berkovitch, A., Deev, K. & Landa, E., 2011. How non-hyperbolic multifocusing improves depth imaging, *First Break,* **29,** 103–111.

Bertalmio, M., Bertozzi, A.L. & Guillermo, S., 2001. Navier-Stokes, fluid dynamics, and image and video inpainting, computer vision and pattern recognition, 2001. CVPR 2001, in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001,* Kauai, HI, USA, 2001, 8–14 December 2001, pp. I–I, doi:10.1109/CVPR.2001.990497.

Goodfellow, I., Bengio, Y. & Courville, A., 2016. *Deep Learning,* MIT Press, http://www.deeplearningbook.org.

Halpert, A., 2018, Deep learning-enabled seismic image enhancement, in *Proceedings of the SEG Annual Meeting,* SEG Technical Program Expanded Abstracts, pp. 2081–2085, doi:10.1190/segam2018-2996943.1.

Hoecht, G., Ricarte, P., Bergler, S. & Landa, E., 2009. Operator-oriented interpolation, *Geophys. Prospect.,* **57,** 957–981.

Kragh, E. & Christie, P., 2002. Seismic repeatability, normalized RMS, and predictability, *Leading Edge,* **21,** 640–647.

Ledig, C. *et al.*, 2017. Photo-realistic single image super-resolution using a generative adversarial network, in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR),* Honolulu, HI, pp. 105–114, doi:10.1109/CVPR.2017.19.

Liu, G., Reda, A.F., Shih, K.J., Wang, T-C., Tao, A. & Catanzaro, B., 2018. Image inpainting for irregular holes using partial convolutions, in *Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science,* Vol. **11215,** eds Ferrari, V., Hebert, M., Sminchisescu, C. & Weiss, Y., Springer.

Mandelli, S., Borra, F., Lipari, V., Bestagini, P., Sarti, A. & Tubaro, S., 2018, Seismic data interpolation through convolutional autoencoder, in *Proceedings of the SEG Annual Meeting,* SEG Technical Program Expanded Abstracts, pp. 4101–4105, doi:10.1190/segam2018-2995428.1.

Ronneberger, O., Fischer, P. & Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation, in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. MICCAI 2015. Lecture Notes in Computer Science,* Vol. **9351,** eds Navab, N., Hornegger, J., Wells, W. & Frangi, A., Springer.

Telea, A., 2004. An image inpainting technique based on the fast marching method, *J. Graph. Tools,* **9**(1), 23–34.

Wang, Y., Wang, B., Tu, N. & Geng, J., 2020. Seismic trace interpolation for irregularly spatial sampled data using convolutional autoencoder, *Geophysics,* **85**(2), V119–V130.

Li, X.R., Mitsakos, N., Lu, P., Xiao, Y. & Zhao, X., 2019. Seismic compressive sensing by generative inpainting network: toward an optimized acquisition survey, *Leading Edge,* **38,** 923–933.

APPENDIX: DNN DESCRIPTION FOR IMAGE INPAINTING

Following Liu *et al.* (2018), we design a Unet-like image inpainting architecture based on partial convolutions, Fig. A1. This DNN contains 16 partial convolutional (PConv) layers, eight upsampling layers and eight concatenation layers (skip connections). The input image and mask dimensions are $512 \times 512 \times 3$ (three channels for each colour component). An activation function for the first eight PConv layers (encoding, or feature extraction, part of the DNN) is ReLU, while the last eight PConv layers (decoding part) have LeakyReLu activation with negative slope coefficient equals to 0.2. The total number of trainable parameters is 32 865 248.

We implement the partial convolution layer by extending the existing convolution layer in TensorFlow. The DNN weights were randomly initialized, and Adam stochastic optimization algorithm was exploited during the training process.

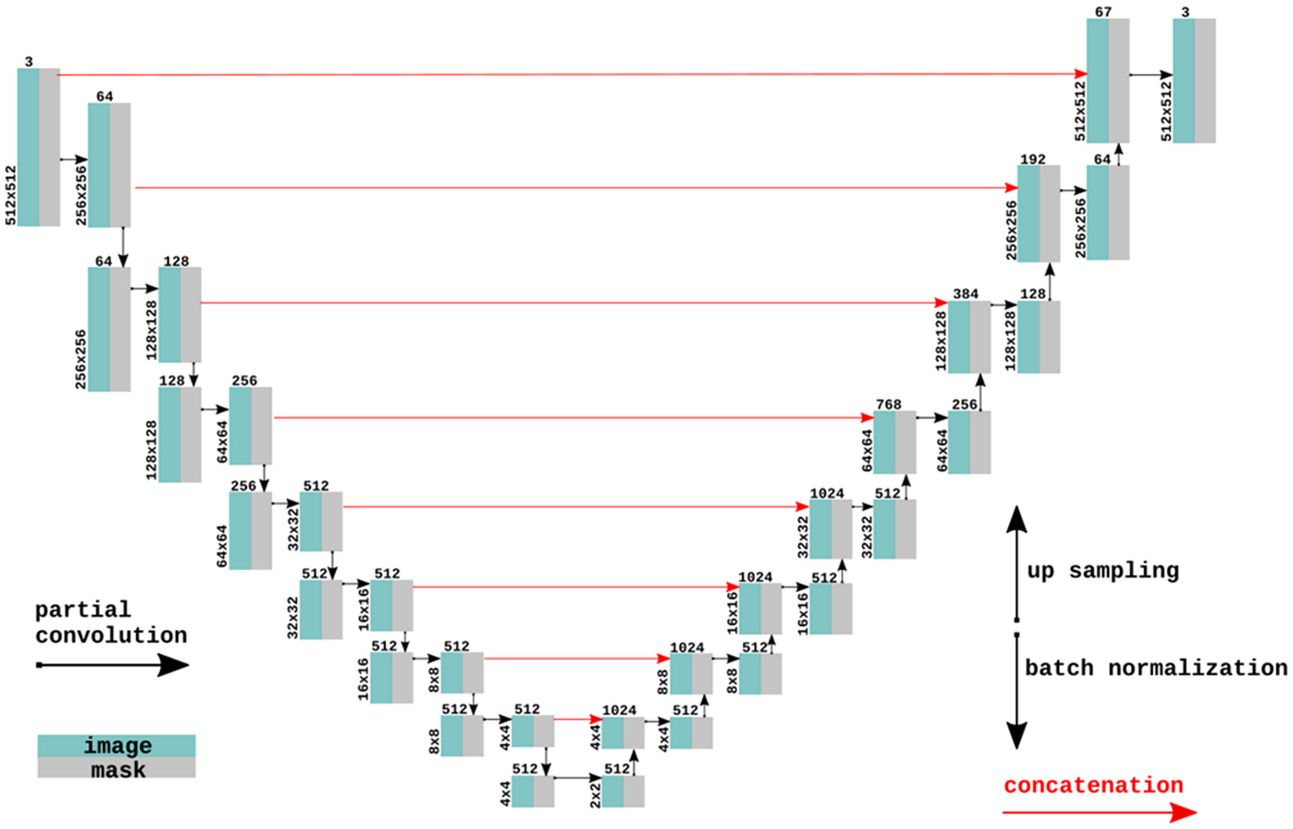

Figure A1. UNet-like DNN architecture for image inpainting with partial convolutions. Arrows on the right label partial convolution layers: upward black arrows indicating upsampling layers, downward black arrows show batch normalization, whereas red arrows denote concatenation layers.