A deep learning-based approach to increase efficiency in the acquisition of ultrasonic non-destructive testing datasets

Nick Luiken and Matteo Ravasi
King Abdullah University of Science and Technology, Thuwal 23955
{nicolaas.luiken, matteo.ravasi}@kaust.edu.sa

Abstract

Ultrasonic phased array systems traditionally acquire data in a sequential fashion. Although using different excitation delays for each pulsing element can be used to steer the emitted wavefield into e.g., plane waves or focused beams, the overall frame rate of the system is dominated by the choice of the firing time between two consecutive experiments. Inspired from a technology in reflection seismology, we propose the use of simultaneous shooting to increase the flexibility of acquiring ultrasonic data in non-destructive testing (NDT) applications. Simultaneous shooting *is an acquisition setup whereby separate transmit sequences are performed simultaneously at a reduced time interval leading to entangled data that may yield artifacts in subsequent imaging products. The data can be untangled by a process called deblending, which is a heavily underdetermined linear inverse problem. We solve the deblending problem using the recently introduced SSDeblend algorithm. This algorithm combines the physics of simultaneous shooting with a powerful self-supervised denoiser specifically tailored to remove the so-called blending noise. We conduct an experiment on an openly available Full Matrix Capture dataset and show that one can speed up the acquisition by at least a factor of 2 with little loss of quality on the resulting final image.*

1. Introduction

Ultrasonic phased array systems have become mainstream in non-destructive testing (NDT) in recent years [7]. Compared to single element transducers, phased arrays are more versatile in that they are capable of simultaneously transmitting and receiving ultrasonic signals across multiple electric transducers. By doing so, NDT operators can perform efficient inspections with tailored ultrasonic beams at various focal lengths and steering angles. Popular configurations are, for example, the plane-wave insonification (PWI) and sector B-scan modes [9]. However, in many areas of industrial NDT the target is static: in such cases, data analysis can be performed off-line and a different acquisition setup that utilises the complete set of time-domain data from all combinations of transmitters and receivers may be preferred. This approach is referred to as full matrix capture (FMC).

A common characteristic of the above mentioned acquisition modalities is represented by the choice of the maximum recording time (T) for a single emission (or its frame rate $f_R = 1/T$). This is usually dictated by distance between the array and the object and the size of the object that we wish to probe and the amount of time required by its multiple reverberations to attenuate below a given noise level, such that the data acquired by the next emission is not polluted by the coda of the previous one. The requirement to wait for the wavefield to attenuate represents a critical factor to keep into account when identifying the overall acquisition time of a NDT acquisition.

Reflection seismology is a remote sensing method that bears strong similarity with the ultrasound acquisition techniques used in NDT. More precisely, seismic data are commonly acquired in a FMC-style mode, with the main difference that a single physical source is available and physically shifted after every firing to acquire data over a line (in 2D) or a grid (in 3D). Receivers may be also fixed like in FMC systems or move alongside the source (towed behind the source within a cable called streamer). The presence of a single source makes it therefore impossible to physically create focused or plane wavefronts, limiting the degrees of freedom to design faster acquisition systems.

To this end, the geophysical community has developed a technology called *simultaneous shooting* with the aim of reducing acquisition time and cost [2, 3, 14]. Simply put, instead of firing a source and waiting for all the reflected energy to bounce back from the subsurface (or at least all of the energy above a given threshold), simultaneous shooting acquisitions fire consecutive sources at shorter time intervals. This leads to the late reflections produced by a given shot to overlap with the early reflections from the consecutive shot. Whilst reducing the acquisition time by a factor

(directly linked to the amount of overlap between consecutive shots), this acquisition modality calls for an extra data processing step, called *deblending*, to produce an estimate of the ideal data that would have been recorded without overlap. The concept of simultaneous shooting has been recently introduced in NDT ultrasound [8, 15], with deblending being approached using standard state-of-the-art geophysical algorithms.

Our contribution

In this work we apply for the first time a deep-learning deblending method to NDT data using the recently proposed deblending algorithm of [13]. This algorithm leverages the power of self-supervised, blind-spot neural networks [4, 11] whilst at the same time incorporating the physics of the blending process in the learning process. We further discuss how this algorithm can be easily adapted to both the FMC and PWI scenarios and perform an in-depth analysis of the cost-quality trade-off introduced when acquiring NDT ultrasound datasets in blended mode.

2. Methodology

In this section, we describe the blending operator that can be used to model blended data from the clean unblended data and, how to cast deblending as an inverse problem, and finally our proposed algorithm, named SSDeblend.

Terminology

Throughout this article, we assume that the recorded radio-frequency data $d(x_s, x_r, t)$ is arranged into a cube of size $n_s \times n_r \times n_t$, where n_s denotes the number of experiments (also loosely referred to as shots), n_r the number of active transduceer elements, and n_t the number of time samples. In FMC, the number of possible transmits coincides with the number of transducer elements, whilst in PWI mode this corresponds to the number of excited plane waves. If we slice the data cube by fixing one shot we refer to this as a *common shot gather* (CSG), whilst if we fix one receiving channel we call the resulting slice a *common receiver gather* (CRG).

Blending

Generally, to reduce the acquisition burden of transmitting on all elements, sparse FMC acquisitions skip a certain number of elements (traditionally transmitting every 2, 4, or 8 elements). Blending, on the other hand, combines recordings from independent transmits by simultaneously performing separate transmit sequences. This process can be accomplished at acquisition time as schematically depicted in Figure 1. More specifically, in FMC mode two consecutive actuators are fired closely spaced in time (Figure 1a), whilst in PWI mode two consecutive plane waves are fired closely spaced in time (Figure 1b).

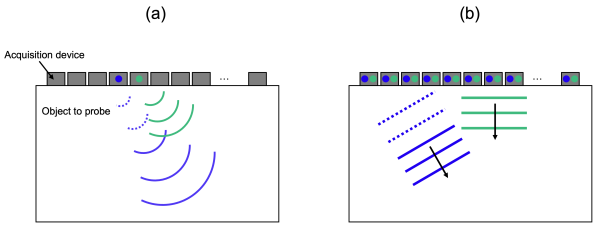

Figure 1. Schematic representation of blended acquisition for a) FMC and b) PWI modalities. Blue waves are emitted first, and green waves are emitted after a given delay time. Small blue and green dots inside the acquisition device indicate which emitters are involved in each experiment.

Numerically speaking, blended data can be described as time shifted unblended data due to the fact the acquired data are simply recorded earlier in time. Every source is fired according to a predetermined schedule called the *dithering code* that relates the clean unblended data to the blended data. This is constructed by choosing a periodic *nominal firing time*, which is then perturbed by a time jitter drawn from a given random distribution. The relation between the clean unblended data and the blended data is given by

$$Bd_c = d_b, \tag{1}$$

where d_c denotes the clean data acquired in a conventional way, d_b is the blended data, and B is the blending operator. Considering the so-called *continuous blending* mode, where a single source is fired at a reduced interval, and hence singly scattered events (primaries) of the consecutive shots overlap the multiply-scattered energy (multiples) of the previous one,continuous blending, we have

$$B = \left[\Phi_1 \cdots \Phi_{n_s}\right],$$

where $\Phi_i \in \mathcal{R}^{(t_{n_s}/\Delta_t + n_t + 1)\times(n_t \cdot n_r)}$ are time shift operators and t_{n_s} denotes the firing time of the last source. Writing $d_{c,i}$ for shot gather i, the blended data are

$$d_b = \Phi_1 d_{c,1} + \ldots + \Phi_{n_s} d_{c,n_s} \tag{2}$$

Equation (1) is an underdetermined problem and therefore has the minimum norm solution

$$d_c = B^H (BB^H)^{-1} d_b = B^H D^{-1} d_b,$$

where the matrix D is a diagonal matrix with D_{ii} corresponding to the number of overlapping shots at that particular time index. Temporarily ignoring the scaling factor D^{-1}, the reconstructed shot gathers are

$$d_{c,i} = \Phi_i^H \Phi_1 d_{c,1} + \ldots + d_{c,i} + \ldots + \Phi_i^H \Phi_{n_s} d_{c,n_s}.$$

The action of the adjoint is referred to as *pseudodeblending*, as it retrieves each shot gather overlapped with the previous

and consecutive shots, which can be considered as *blending noise*. Since the blending noise has the same structure as the signal, it will be extremely hard to filter it. Luckily, when re-sorting pseudodeblended data into CRGs, the blending noise appears as trace-wise coherent noise that can be filtered by finding an appropriate domain in which the signal and noise can be discriminated. The "randomness" of the noise depends on the jitter of the dithering code, as larger jitters will lead to a bigger difference in firing time between consecutive shots. Conversely, very small dithers will yield almost coherent events making them harder to filter out. Figure 2 illustrates the unblended and the pseudo-deblended CSG and CRG for the dataset used in the Numerical Results section. Deblending may be cast as an inverse problem and

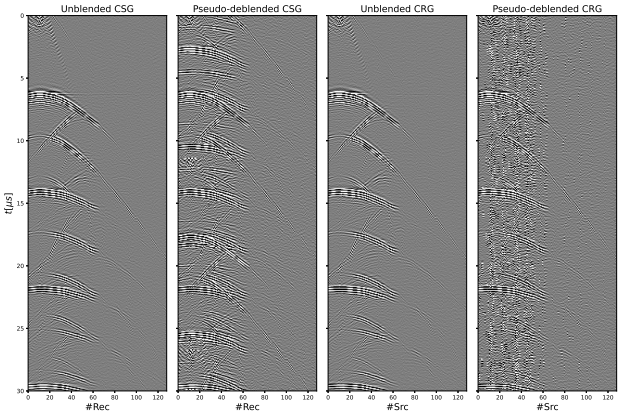

Figure 2. Unblended and pseudo-deblended CSG and CRG.

can be solved by the following optimization problem:

$$x_\star = \arg\min_x \frac{1}{2}\|BPFx - d_b\|_2^2 + \mathcal{R}(x),\ d_c = PFx_\star \quad (3)$$

Here, $\mathcal{R}(\cdot)$ is typically a norm, F is a transform to some suitable domain where the regularization applies. P acts as a transform to the data to which F applies, which can be an operator that extracts patches from x or requires F to apply to CSGs or CRGs. Common approaches are to use $\mathcal{R}(\cdot) = \lambda\|\cdot\|_1$ where F transforms to either the frequency-wavenumber [1] or the Curvelet [12] domain, or $\mathcal{R}(\cdot) = \lambda\|\cdot\|_\star$ (the nuclear norm) [10].

Self-supervised denoising

Key to a successful deblending algorithm is the availability of an appropriate filter for the trace-wise blending noise. Inspired by recent advances in machine learning, we propose the use of a *self-supervised denoiser*. Self-supervised denoisers are neural networks that aim to denoise data without access to clean labels for training. Instead, the input is masked and the label is the unmasked noisy data. This way, the network is forced to infer the data in the masked region by using information from neighbouring pixels, since the network never sees the pixel itself. Because the pseudo-deblended CRGs contain trace-wise structured noise, it makes sense to mask a trace and let the network infer the underlying signal from the neighbouring traces. This approach is called StructNoise2Void, and was introduced by [4]. Our specific implementation of this network is slightly different in that we construct a network with a so-called blind spot by creating a receptive field that only sees the top part of the image through padding and cropping the input appropriately. This is an adaptation of the network introduced by [11], who used this approach to create a blind pixel network for random noise suppression. The network, which we call StructBS, is illustrated in Figure 3. The benefit of StructBS is that it does not require any clean labels, which in the case of deblending are commonly not available.

Plug-and-Play regularization

Although StructBS is a powerful denoiser for the pseudo-deblended data, at first sight it seems that it can only be used as a denoiser. Ideally, we would like to use StructBS also in an inversion approach; however, it is not straightforward how to incorporate such denoiser into an inversion scheme since it cannot be added as a regularization term as in equation 3. [16] introduced a methodology derived from the Alternating Direction Method of Multipliers (ADMM) called Plug-and-Play (PnP) regularization that allows to incorporate any denoiser into an inversion scheme. For completeness, we give a short derivation of the method in the context of the deblending problem, but it should be understood that both the ADMM and the PnP algorithm apply to more general problems, even nonlinear ones.

The ADMM algorithm is generally applied to inverse problems of the form in equation 3 when $\mathcal{R}$ is non-smooth. In this case one cannot apply standard gradient-based optimization methods because the resulting objective function is non-smooth. ADMM resolves this issue by applying a variable splitting strategy (i.e., $x = y$), and solving

$$\min_{x,y} \frac{1}{2}\|Bx - d_b\|_2^2 + \mathcal{R}(y) \text{ subject to } x = y$$

by forming the so-called *augmented Lagrangian*,

$$\max_u \min_{x,y} \frac{1}{2}\|Bx - d_b\|_2^2 + \mathcal{R}(y) + \frac{\rho}{2}\|x - y\|_2^2 + u^T(x - y),$$

where u is the Lagrange multiplier and ρ is a scalar. By alternatively optimizing over the variables x, y and u we obtain the ADMM iterations:

$$\begin{aligned} y_{k+1} &= \arg\min_y \left\{\mathcal{R}(y) + \frac{\rho}{2}\|x_k - y + u_k\|_2^2\right\} \\ x_{k+1} &= \arg\min_x \left\{\frac{1}{2}\|Bx - d_b\|_2^2 + \frac{\rho}{2}\|x - y_{k+1} + u_k\|_2^2\right\} \\ u_{k+1} &= u_k + x_{k+1} - y_{k+1}. \end{aligned}$$

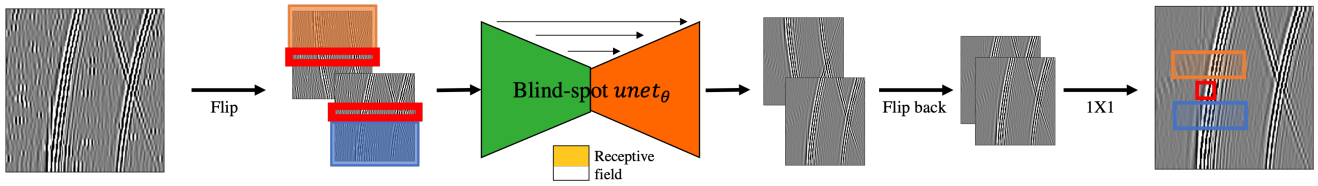

Figure 3. Blind-spot network. By padding and accordingly cropping the image a receptive field is created that is blind to a specific trace. The receptive field is indicated by the orange and blue color in the input and the blind trace is indicated by the red rectangle.

The minimization for the y-update is called the *proximal operator*. For the most widely used regularization terms, like $\mathcal{R}(\cdot) = \lambda\|\cdot\|_1$ and $\mathcal{R}(\cdot) = \lambda\|\cdot\|_\star$, the proximal operator has a closed form solution that is cheap to apply. [16] observed that the proximal operator of a function $\mathcal{R}$ is equivalent to a denoising problem with regularization $\mathcal{R}$. Given the abundance of advanced denoisers that can not be cast in this particular regularized form, they proposed to drop this formulation and replace the y-update with a denoiser of choice. Examples include non-local means [5], BM3D [6] or denoisers based on neural network architectures (e.g. [17–19]).

SSDeblend

In [13] the authors introduce the *SSDeblend* algorithm, which we shortly recap here. Following the PnP methodology, the authors propose the use of a self-supervised denoiser, named StructBS, to filter the trace-wise coherent noise in CRGs within an inversion workflow. This leads to the following algorithm:

$$\begin{aligned} y_{k+1} &= \text{StructBS}_\theta(x_k + u_k) \\ x_{k+1} &= \arg\min_x \left\{ \frac{1}{2}\|Bx - d_b\|_2^2 + \frac{\rho}{2}\|x - y_{k+1} + u_k\|_2^2 \right\} \\ u_{k+1} &= u_k + x_{k+1} - y_{k+1}. \end{aligned}$$

The x-update is solved by means of LSQR, which for deblending only requires a few iterations, and the StructBS network is trained on-the-fly and updated at every iteration. Finally, at the end of the iterations, we apply a post-processing step that is meant to minimize the signal leakage and involves solving the following optimization problem:

$$d_c = d_{\text{SSDeblend}} + \arg\min_x \frac{1}{2}\|B(x + d_{\text{SSDeblend}}) - d_b\|_2^2.$$

where $d_{\text{SSDeblend}}$ is the final x-update. This post-processing is meant to force the data residual to be exactly 0, which means that all the amplitudes in the deblended data match the amplitudes of the true data. Due to the nature of the blending operator this will happen only in areas where there is no overlap between consecutive shots. The post-processing step is also solved by means of LSQR: in this case we choose 5 steps.

Imaging modalities

In this section, we briefly present the different imaging modalities used in the Numerical Results section. Whilst comparing the wavefields before and after deblending with the ideal unblended data can already provide us with an indication of how data quality deteriorates when acquiring data in simultaneous shooting mode, it is also important to asses how such blending noise translate into potential imaging artefacts, given that any further manual or automated interpretation is commonly done on the resulting images.

The total focusing method (TFM) is the most commonly used imaging algorithm for FMC datasets. This algorithm constructs an image by a (possibly weighted) summation of all data contributions that may have originated from a given image point. In longitudinal-longitudinal (LL) imaging, given a longitudinal velocity v_L, the total traveltime for any emitter-receiver-image point triplet is defined as:

$$\begin{aligned} t_{s-i-r} &= t_{s-i} + t_{i-r} \\ &= \frac{\sqrt{(x_s - x_i)^2 + z_i^2}}{v_L} + \frac{\sqrt{(x_r - x_i)^2 + z_i^2}}{v_L} \end{aligned} \tag{4}$$

and the image is constructed for each image point:

$$I_{LL}(x_i) = \sum_{s=1}^{n_s}\sum_{r=1}^{n_r} d(x_s, x_r, t_{s-i-r}) \tag{5}$$

Equation 4 must be modified to perform PW imaging. In this scenario, whilst the receiver-side traveltime remains untouched, the source-side traveltime is modified to account for propagation parallel to the direction of the emitted plane wave:

$$\begin{aligned} t_{p-i-r} &= t_{p-i} + t_{i-r} \\ &= \frac{(x_i - x_0)\sin\theta + z_i\cos\theta}{v_L} + \frac{\sqrt{(x_r - x_i)^2 + z_i^2}}{v_L} \end{aligned} \tag{6}$$

where x_0 is the horizontal location of the first (for positive angles θ) or last (for negative angles θ) receiver. Note that the numerator of t_{p-i} can be simply interpreted as the distance between the image point x_i and a straight line passing through x_0 with angle from the horizontal axis equal to θ.

3. Numerical Results

Dataset and data pre-processing The proposed algorithm is tested on a openly available FMC dataset, acquired at the University of Strathclyde (Glasgow) [1]. The dataset was acquired to inspect a welded stainless steel plate with longitudinal wave speed equal to $v_L = 5820m/s$ and thickness of 22mm; the sample contains a 6mm lack-of-fusion flaw at an angle of 50 degrees with respect to the x-axis. The acquisition device is a typical 5 MHz ultrasonic phased array with 128 elements with an element pitch of 0.7 mm and sampling frequency of 100 MHz. For each emission, the total recording time is $100\mu s$.

To begin with, a basic pre-processing sequence is implemented to mitigate acquisition noise in the data: first a trapezoidal filter with cut-off frequencies of [2, 2.5, 7.5, 8] MHz is applied to the data, followed by resampling at 16 MHz. Finally, the data is filtered in the frequency-wavenumber domain around the zero horizontal wavenumber ($k_x = 0$) to suppress horizontally planar events likely due to the probe construction. The original and pre-processed data are shown in Figure 4. As a final step we truncate the data to 1024 time samples (total time duration of $64\mu s$) in order to have a compatible size for the UNet architecture. Finally, a PW dataset in synthesised from the physically recorded FMC data. This second dataset mimics a scenario with 64 plane waves with angles regularly sampled from -45 to 45 degrees (Figure 5).

Deblending To test the performance of SSDeblend we consider a number of different dithering codes, where both the nominal firing time and the random jitter are varied. Given that the firing interval for the conventionally acquired FMC data is roughly 64 μs (after truncation in pre-processing), we choose nominal firing times of 32 μs, 16 μs, and 6.4 μs, corresponding to an overlap of 50%, 75% and 90% between two consecutive shots. Clearly, more overlap leads to faster acquisition at the cost of a more challenging deblending, and therefore deblended data with reduced signal-to-noise ratio (SNR). We construct random jitter by drawing from a normal distribution with standard deviations of 10μs (for 32 μs and 16 μs nominal firing time), 5μs, and 1μs (for all cases). All experiments are carried out on a NVIDIA Tesla V100 GPU.

A reconstructed CRG and CSG are first shown in Figure 6 for the dithering code 32 μs $\pm$ 10 μs and the FMC data and the PW data, respectively. Next, we analyse the trade-off between the reduction in acquisition time versus the SNR of the deblended data for the FMC data set (Figure 7). Clearly, the choice of dithering code depends on the balance between the loss of data quality that is acceptable and the necessity for fast acquisition.

[1] DOI:10.15129/a7da5071-0436-4913-8f0e-de78f5ebccd6

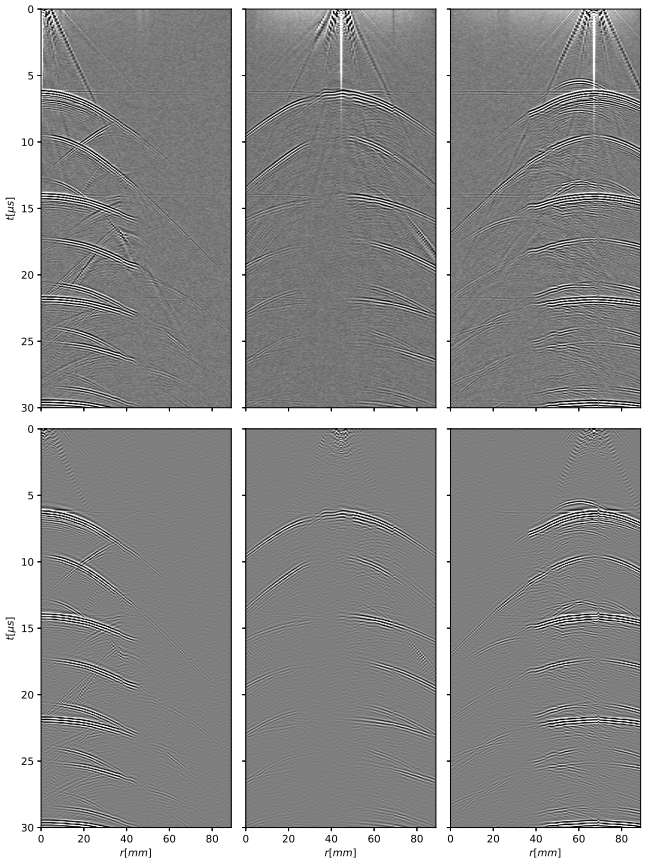

Figure 4. Raw and pre-processed FMC data. The top panel shows some raw CSGs and the bottom panel the corresponding preprocessed CSGs.

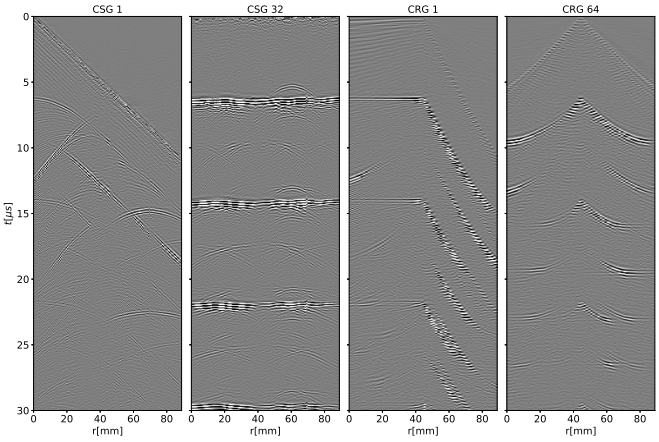

Figure 5. Two CSGs and CRGs showing the synthesized plane wave dataset.

Imaging Finally, to further assess the quality of our deblending algorithm we image the deblended data and compare the resulting images with those from the unblended and

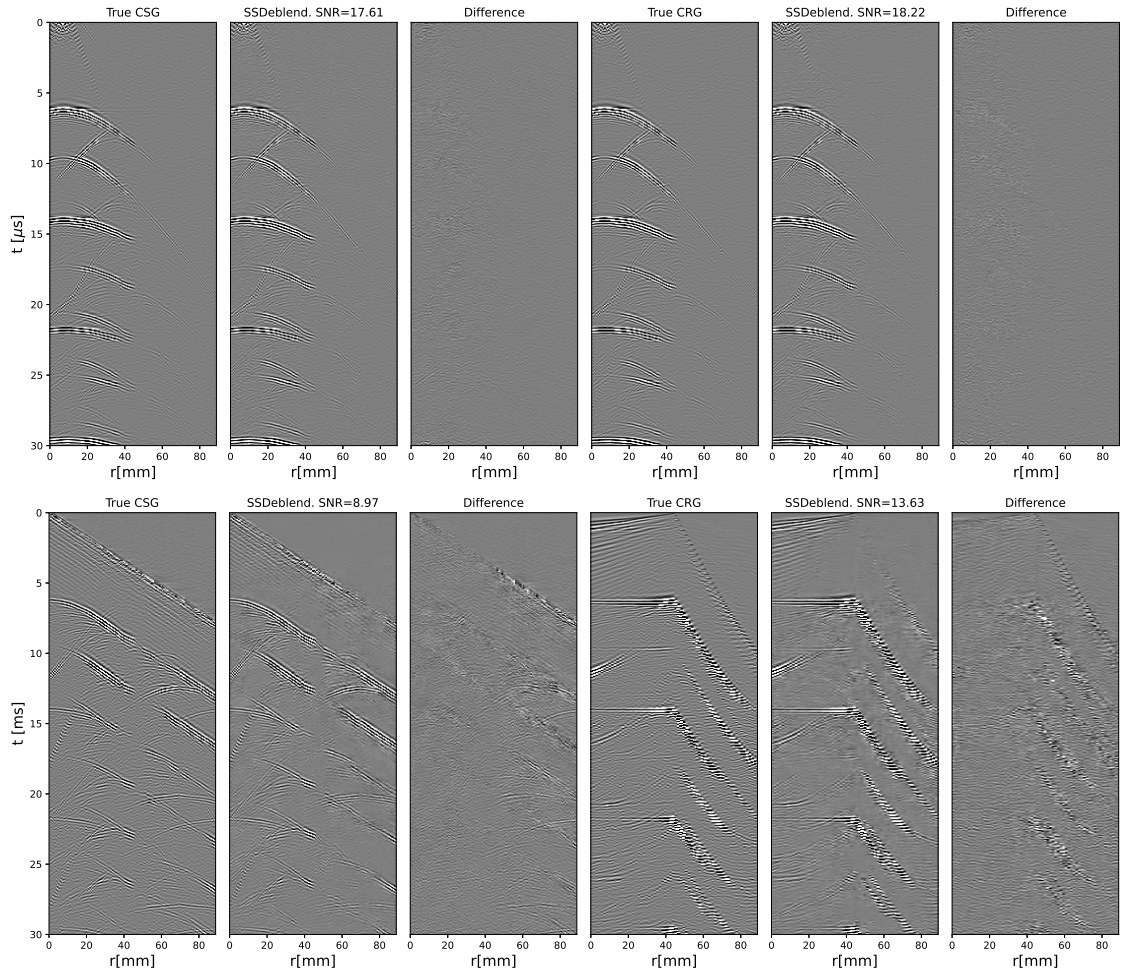

Figure 6. True data, deblended data and difference for a CSG and a CRG with a nominal firing time of 32μs and a dither of 5μs. The top panel shows the FMC data and the bottom panel the plane wave data.

the corresponding pseudo-deblended data. First, we show the migrated data from the unblended data in Figure 8(a). Figure 8(b) shows that the image formed directly from the pseudo-deblended data is significantly polluted by noise, making it difficult to detect the flaw, especially for the plane wave data. In this case, both the best and worst deblending result for the FMC data still capture the flaw quite well, and the noise introduced by imperfect deblending is mainly of random nature and therefore does not distort the image. Interestingly, the images in Figure 8(c) for the FMC data seems to have less noise than those from the original unblended data, especially on the top left and right below the gap in the steel plate. Moreover, the wave-like artifact is less pronounced in Figure 8(c), and completely absent in Figure 8(d). This may be an indication that our deblending algorithm acts also as a denoiser of incoherent features present in the data, which are unrelated to the blending process.

In general, the images from the PW datasets are of lower quality to their FMC counterparts. For example, the flaw in the image from unblended PW data looks smaller than in the FMC image. The bottom panel of Figure 8(c) shows that the best deblending result captures the flaw quite well. On the other hand, the image from worst deblending results seems quite deteriorated, although the flaw is still clearly visible. We conclude that for this particular case, one would prefer the fastest acquisition since the anomaly of interest is still imaged well, but the anomaly is also quite simple. For more complicated anomalies, deblending scenarios that lead to a higher SNR may be required.

Ablation SSDeblend depends on a number of hyperparameters. Firstly, we address the number of iterations for LSQR. The convergence rate of LSQR depends on the condition number and the clusters of eigenvalues. The eigenvalues are determined by the number of overlapping shots at a particular point, and hence the largest eigenvalue of

the blending operator is the largest number of overlapping shots, and the smallest non-zero eigenvalue is 1. Since the condition number is low and there is a large number of repeated eigenvalues, LSQR converges in very few steps. Therefore, we choose 3 iterations of LSQR. The other hyperparameters are related to the StructBS denoiser. The first important parameter is the architecture of the network itself. In our work we have used the architecture of [11], which is a 5-layer UNet with 48 channels in the first layer, a convolutional filter of size 3×3, and a leaky ReLU activation function with a slope of -0.1. Training is performed using the Adam optimizer, whose learning rate is also a hyperparameter, along with the batch size of the data and the number of training epochs. We show the SNR as a function of outer iterations for the FMC dataset with a nominal firing time of 16 μs and a dither of 10 μs in Figure 9 for different batch sizes and learning rates.

With regard to the number of denoising epochs we have found that less than 40 epochs consistently leads to a lower SNR after 40 outer iterations. However, after more outer iterations the algorithm may be able to catch up when the same number of total denoising epochs is reached, i.e. number of epochs $\times$ outer iterations.

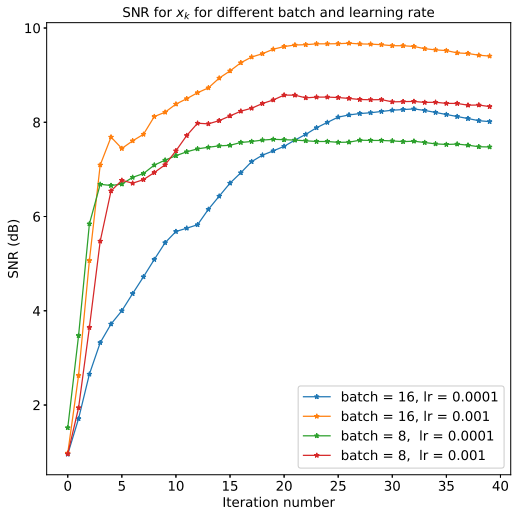

Figure 9. SNR as a function of outer iterations for different batch size and learning rate.

4. Limitations

One possible drawback of the proposed method is the computational cost associated with the deblending process. Although the acquisition time may be greatly reduced, the subsequent processing of the data may represent a significant burden for real-time applications. Trading acquisition time for processing time is instead very appealing for offline applications. To mitigate such a limitation, one could potentially reduce the number of iterations to a given preselected computational time, thereby accepting a loss in SNR; overall, we envision deblending to be an integrated part of the process of designing a NDT acquisition system.

Another scenario where the proposed algorithm may struggle with is the case of decimated transmit sequences, where only part of the transmitters (or few plane waves) are emitted to speed up acquisition. In this case, CRGs will be aliased making it difficult for our blind-spot network to leverage signal coherency in nearby traces to predict the blind trace. However, as simultaneous shooting already reduces the acquisition time, the need for decimated transmit sequences may be offset by the reduced transmit interval between subsequent shots.

5. Conclusion

We have proposed the use of simultaneous shooting to speed up the acquisition process of ultrasonic phased array systems. Conventional acquisition systems rely on large time delays between consecutive shots in order for the receivers to record all of the reflected energy originated by a given shot. In simultaneous shooting mode, shots are fired at a reduced time interval thereby decreasing the acquisition time; however, as the recorded shot gathers overlap, the resulting images may be polluted by unwanted artifacts. Therefore, a deblending step is required to disentangles the overlapping shot gathers yielding (at least, in case of perfect deblending) the data that would have been acquired by conventional acquisition.

Deblending is here achieved by means of SSDeblend, an algorithm that combines the physics of simultaneous shooting with a highly effective self-supervised denoiser. Being self-supervised, such denoiser does not require any clean labels, pre-training, or domain adaptation to data that are not seen during training. We have tested SSDeblend on an openly available FMC dataset and a synthesised PW dataset. Both data are blended artificially with different dithering codes corresponding to a roughly 50%, 75% and 90% reduction in acquisition time. Our experiments show that SSDeblend is able to deblend NDT data effectively, and that the quality of the reconstruction is inversely proportional to the reduction in acquisition time, leading to a trade-off between data quality and acquisition time. Based on a visual inspection of the resulting images, we argue that for this particual case the acquisition time can be reduced by 90% while still being able to detect the anomaly in the image for the FMC data. For the PW data, the image for a 90% reduction in acquisition time is less clear, but the anomaly can still be detected. Whilst for this particular experiment the flaw was rather simple to detect, we suspect that in more challenging scenarios a 90% reduction in acquisition time is unrealistic, whereas 50% or even 75% still seems reasonable.

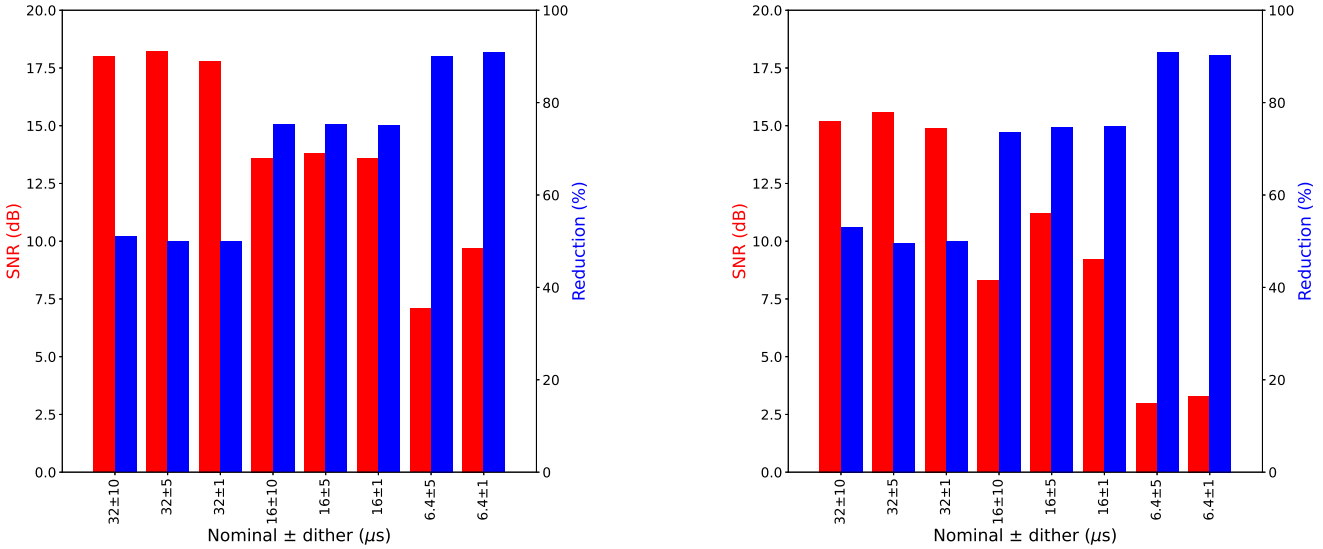

(a) FMC dataset.

(b) Plane wave dataset.

Figure 7. Comparison of the SNR versus the reduction in acquisition time for different nominal firing times and dithers for the plane wave dataset. Notice the clear inverse relation between the two variables. The SNR for the plane wave dataset is considerably lower than for the FMC dataset.

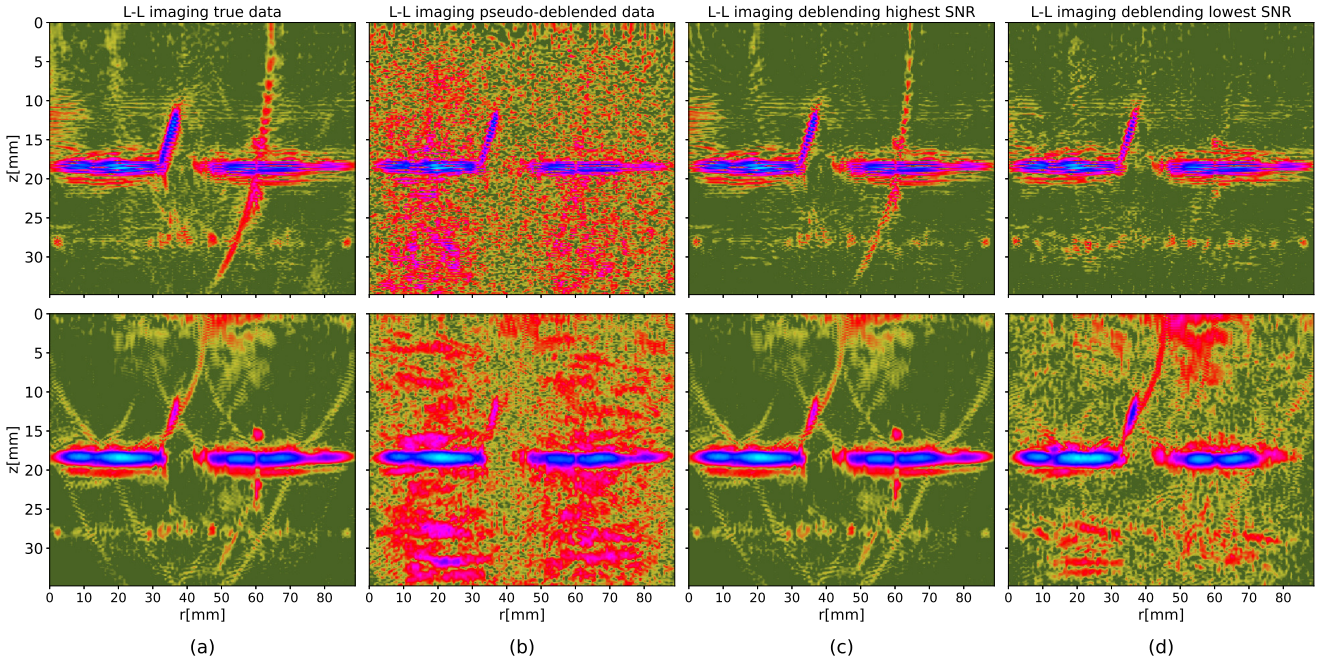

Figure 8. Results of L-L migration for the true data, the pseudo-deblended data and the debleding with the highest and lowest SNR. The top panel corresponds to the FMC data and the bottom panel to the plane wave data.

References

[1] R. Abma, D. Howe, M. Foster, I. Ahmed, M. Tanis, Q. Zhang, A. Arogunmati, and G. Alexander. Independent simultaneous source acquisition and processing. *Geophysics*, 80(6):WD37–WD44, 2015. 3

[2] C. Bagaini. Acquisition and processing of simultaneous vibroseis data. *Geophysical Prospecting*, 58:81–99, 2010. 1

[3] A. J. G. Berkhout. Changing the mindset in seismic data acquisition. *The Leading Edge*, 27:924–938, 2008. 1

[4] C. Broaddus, A. Krull, M. Weigert, U. Schmidt, and G. Myers. Removing structured noise with self-supervised blind-spot networks. *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 159–163, 2020. 2, 3

[5] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. *IEEE Conference on Computer Vision and Pattern Recognition*, 2:60–65, 2005. 4

[6] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 4

[7] Bruce W. Drinkwater and Paul D. Wilcox. Ultrasonic arrays for non- e evaluation: A review. *NDT&E International*, 39:25–541, 2006. 1

[8] Arno Duijster, Gert-Jan van Groenestijn, Paul van Neer, Gerrit Blacquière, and Arno Volker. Blending of phased array data. *AIP Conference Proceedings*, 1949(1):080005, 2018. 2

[9] Caroline Holmes, Bruce W. Drinkwater, and Paul D. Wilcox. Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation. *NDT&E International*, 38:701–711, 2005. 1

[10] R. Kumar, H. Wason, and F.J. Herrmann. Source separation for simultaneous towed-streamer marine acquisition — a compressed sensing approach. 6:WD73–WD88, 2015. 3

[11] S. Laine, T. Karras, J. Lehtinen, and T. Aila. High-quality self-supervised deep image denoising. *Advances in Neural Information Processing Systems*, pages 6970–6980, 2019. 2, 3, 7

[12] C. Li, C.C. Mosher, and Y. Ji. Randomized marine acquisition with compressive sampling matrices. *Geophysical Prospecting*, 60(3):648–662, 2019. 3

[13] Nick Luiken, Matteo Ravasi, and Claire Birnie. A hybrid approach to seismic deblending: when physics meets self-supervision. *arXiv*, 2022. 2, 4

[14] I. Moore, B. Dragoset, T. Ommundsen, D. Wilson, C. Ward, and D. Eke. Simultaneous source separation using dithered sources. *SEG Technical Program Expanded Abstracts*, 2008. 1

[15] Elodie Tiran, Thomas Deffieux, Mafalda Correia, David Maresca, Bruno-Felix Osmanski, Lim-Anna Sieu, Antoine Bergel, Ivan Cohen, Mathieu Pernot, and Mickael Tanter. Multiplane wave imaging increases signal-to-noise ratio in ultrafast ultrasound imaging. *Physics in Medicine and Biology*, 60(21), 2015. 2

[16] S.V Venkatakrishnan, C.A. Bouman, and B. Wohlberg. Plug-and-play priors for model based reconstruction. *2013 IEEE Global Conference on Signal and Information Processing*, 2013. 3, 4

[17] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. van Gool, and R. Timofte. Plug-and-play image restoration with deep denoiser prior. In *IEEE TPAMI*, 2021. 4

[18] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, pages 3142–3155, 2017. 4

[19] K. Zhang, W. Zuo, S. Gu, and L. Zhang. Learning deep cnn denoiser prior for image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3929–3938, 2017. 4