Deep Unsupervised 4-D Seismic 3-D Time-Shift Estimation With Convolutional Neural Networks

Dramsch, Jesper Sören; Christensen, Anders Nymark; Macbeth, Colin; Lüthje, Mikael

Published in:
IEEE Transactions on Geoscience and Remote Sensing

Link to article, DOI:
[10.1109/TGRS.2021.3081516](https://doi.org/10.1109/TGRS.2021.3081516)

Publication date:
2021

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Dramsch, J. S., Christensen, A. N., Macbeth, C., & Lüthje, M. (2021). Deep Unsupervised 4-D Seismic 3-D Time-Shift Estimation With Convolutional Neural Networks. *IEEE Transactions on Geoscience and Remote Sensing*, 60. <https://doi.org/10.1109/TGRS.2021.3081516>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Deep Unsupervised 4-D Seismic 3-D Time-Shift Estimation With Convolutional Neural Networks

Jesper Sören Dramsch¹, Anders Nymark Christensen, Colin MacBeth, and Mikael Luthje

Abstract—We present a novel 3-D warping technique for the estimation of 4-D seismic time-shift. This unsupervised method provides a diffeomorphic 3-D time shift field that includes uncertainties, therefore, it does not need prior time-shift data to be trained. This results in a widely applicable method in time-lapse seismic data analysis that is not implicitly biased by supervised time-shifts from other methods. We explore the generalization of the method to unseen data both in the same geological setting and in a different field, where the generalization error stays constant and within an acceptable range across test cases. We further explore upsampling of the warp field from a smaller network to decrease computational cost and see some deterioration of the warp field quality as a result. This method provides an accurate 3-D seismic registration method, where the heavy computation can be preexecuted and the inference of the network taking seconds on consumer hardware.

Index Terms—3-D time-shift, 4-D seismic, deep learning, neural network, time-lapse, unsupervised learning.

I. INTRODUCTION

SEISMIC time-lapse data consist of two 3-D reflection amplitude cubes that represent the subsurface they were collected from. These cubes are acquired years apart with expected changes in the subsurface due to e.g., hydrocarbon production. The differences in the subsurface cause changes in both amplitudes and velocities, which introduces misalignment of seismic reflectors. Measuring the misalignment and aligning these surfaces to obtain a reliable difference cube is one of the main disciplines in 4-D seismic processing.

These time shifts are most commonly obtained by windowed cross correlation and other statistical or signal processing approaches [1]. Considering the recent advances of machine learning in imaging and domain transfer, we explore possibilities of alignment with convolutional neural networks (CNNs). Machine learning approaches, however, most commonly require labeled data to find a mapping $f(x) = y$, with x being the input data, f being the blackbox algorithm like a neural network, and y being the labels or target.

Manuscript received August 14, 2020; revised February 14, 2021; accepted April 18, 2021. This work was supported by the Danish Hydrocarbon Research and Technology Centre through the Advanced Water Flooding Program. (Corresponding author: Jesper Sören Dramsch.)

Jesper Sören Dramsch and Mikael Luthje are with the Danish Hydrocarbon Research and Technology Centre, Technical University of Denmark, 2800 Lyngby, Denmark (e-mail: jesper@dramsch.net).

Anders Nymark Christensen is with the Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Lyngby, Denmark.

Colin MacBeth is with the School of Energy, Geoscience, Infrastructure and Society, Heriot Watt University, Edinburgh EH14 4AS, U.K.

Digital Object Identifier 10.1109/TGRS.2021.3081516

Machine learning in geoscience has a long history of codevelopments [2]. In recent publications that focus on deep neural networks in geoscience, these methods were applied to a variety of seismic applications including seismic inversion [3]–[5], seismic processing [6]–[8], seismic compression [9], as well as, automated seismic interpretation [10]–[12].

A common problem in machine learning for subsurface science is determining the ground truth. Obtaining information from the subsurface is often prohibited by cost, and e.g., core samples are highly localized data that are often altered by the extraction method as well as the sheer act of unearthing the sample. Additionally, synthetic data may introduce the inverse crime [13] of using the same theory to generate and invert data. Luckily, the physics of medical imaging and inversion is very similar to geophysics, where methods can be validated and fine-tuned. The main method discussed in this article is adapted from the medical imaging literature.

The lack of ground truth leads to another problem that deep learning address but do not solve. For classic neural networks, we need to know a target label dataset, i.e., knowing a prior warp velocity. In 4-D seismic this would mean employing an established method to obtain time shifts. This would effectively result in abstracting that method in a neural network, or modeling the warp, which would lead to committing the inverse crime. This leads us to explore unsupervised methods.

We discuss several options for architectures for mapping the monitor seismic cube to the base seismic cube directly within the network. This is possible in unsupervised configurations but depending on the architecture of the network this problem can be ill-constrained and generate nonphysical mappings. One warranted criticism of deep learning and neural networks is the lack of explainability and limited interpretability. However, we employ a deep neural network to obtain warp velocity vectors, a 3-D equivalent of time shifts, for dense deterministic warping instead of directly obtaining the warped result from a neural network. This enables us to interpret the warping vectors and constrain the warp path.

Moreover, we present the first 4-D seismic 3-D time shift estimator with uncertainty measures. We achieve this by implementing a variational layer that samples from a Gaussian with the reparameterization trick [14]. Therefore, we can counteract some of the noise on the performance of the network.

II. THEORY

Extracting time shifts from 4-D seismic data is most commonly done trace-wise (1-D), which limits the problem to depth. This provides sufficient results for simple problems.

However, geologically complex systems and prestack time shifts benefit from obtaining 3-D time shifts. We discuss classical 3-D time-shift extraction methods, we then go on to discuss relevant deep learning methods. These methods extract time shifts with different constraints which we explore. For brevity, we present the results of the best method to date, developed for the medical domain: VoxelMorph [15].

The goal of both conventional and machine learning methods is to obtain a warp velocity field $\mathbf{u}(x, y, z)$ that ideally aligns two 3-D cubes B and M within given constraints. That means a sample $m[x, y, z]$ will be aligned by adjusting $m[x+u_x, y+u_y, z+u_z]$. In image processing, this is considered “dense alignment” or “dense warping,” hence, we need a dense vector field to align each sample in the base and the monitor cube. Generally, $\mathbf{u}(x, y, z) \in \mathbb{R}^3$, which implies interpolation to obtain the warped result.

A. Conventional Methods

Most conventional methods in 4-D seismic warping focus on 1-D methods [16], which include local 1-D cross correlation, dynamic time warping (DTW) [17], local similarity attributes [18], optical flow methods and methods based on Taylor expansion [19]. In total there are over 30 methods [20], we do not cover these methods in detail, but focus on the limited applications of 3-D methods in 4-D seismic warping.

1) *Local 3-D Cross Correlation*: Hall *et al.* [21] introduced local 3-D cross correlation as a method for surface-based image alignment. The horizon-based nodal cross correlation results were then linearly interpolated to full cubes. Hale [22] extended this method to full seismic cubes by calculating the multidimensional cross correlation windowed by a Gaussian with a specified radius. The correlation results are normalized to avoid spurious correlations by amplitude fluctuations and high-amplitude events. Subsequently, the cross correlation result is searched for peaks using the following triple sum:

$$c[u_x, u_y, u_z] = \sum_{x, y, z=-\infty}^{\infty} b[x, y, z] \cdot m[x + u_x, y + u_y, z + u_z] \quad (1)$$

with c being the cross correlation lag. The computational complexity of this method is $\mathcal{O}(N_s \times N_l)$ with N_s being the total number of samples and N_l being the total number of lags.

Stabilization of the results of 3-D cross correlation is obtained by applying spectral whitening of the signals and smoothing the images with a Gaussian filter without increasing the computational complexity despite the windowing function [22].

2) *Inversion-Based Methods*: Rickett *et al.* [23] describe a nonlinear inversion approach, with the objective function being

$$\mathbb{E} = |\mathbf{d} - f(\mathbf{m})|^2 + |\nabla_x(\mathbf{m})|^2 + |\nabla_y(\mathbf{m})|^2 + |\nabla_z^2(\mathbf{m})|^2 \quad (2)$$

with $\mathbf{m}$ being the model vector, $\mathbf{d}$ being the data vector. The nonlinear inversion is constrained by applying the first derivative to the spatial dimensions x , y , and Laplacian in z to obtain a smooth solution. Cherrett *et al.* [24] implement a geostatistical joint inversion that uses the geostatistical

information combined with data constraints as a prior in a Bayesian inversion scheme

$$P(\mathbf{x}|\text{geostats}, \text{data}) \propto \exp(-(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu})/2) \quad (3)$$

with $\mathbf{C}$ being the posterior covariance matrix, $\mathbf{x}$ the sample mean vector, and $\boldsymbol{\mu}$ being the posterior mean vector.

3) *Medical Imaging*: According to [25], the rich history of medical image registration consists of four main steps, such as feature detection, feature matching, transform model estimation, and image resampling and transformation. Within the scope of this article, transform model estimation is the main interest, which defines a mapping function from the base image to the moving image. The transformation models fall into several general categories. Global mapping models define a global transformation of the entire image, which is unsuitable to this application of 4-D seismic. Local mapping models have been shown to outperform global methods [25] and include piecewise mappings and weighted least squares [26]. Alternatively, transforming the moving image through radial basis functions and matching a globally linear model matches images with significant local distortion [25]. Finally, elastic matching presents a nonrigid registration method [27] that finds an optimal matching between images according to intensity values and boundary conditions such as smoothness and stiffness of the matching vectors [28]. Diffeomorphic mapping is not explicitly outlined in [25], but particularly relevant to this article. In [29], large deformation flows were put forth that greedily find a path through diffeomorphic transformations. Diffeomorphisms have gained great attention in the medical field, particularly with large deformation diffeomorphic metric mapping (LDDMM) [30]. This method iteratively finds the shortest path through small diffeomorphisms and is computationally expensive, which is a possible explanation that they have not found greater use in geophysics, due to larger datasets.

B. Machine Learning Methods

The machine learning methods discussed in this section are imaging based, and therefore, rely on recent advances of CNNs in deep learning. We discuss different approaches that include supervised and unsupervised/self-supervised methods. These methods are all based on CNNs.

CNNs are a type of neural network that is particularly suited to imaging approaches. They learn arbitrary data-dependent filters that are optimized based on the chosen objective via gradient descent. These filters can operate on real images, medical images, or seismic data alike. The convolutional filter benefits from weight sharing, making the operation efficient and particularly suited to GPUs or specialized hardware. In Fig. 1, we show a schematic, that is convolved with moving 3×3 filters repeatedly to obtain a spatially downsampled representation. These convolutional layers in neural networks can be arranged in different architectures that we explore in the following analysis of prior methods in image alignment.

1) *Supervised CNNs*: Supervised end-to-end CNNs rely on reliable ground truth, including the time shifts being available. Training a supervised machine learning system requires both a

Fig. 1. Schematic CNN. The input layer (yellow) is convolved with a 3×3 filter that results in a spatially subsampled subsequent layer that contains the filter responses. This second layer is again convolved with a 3×3 filter to obtain the next layer. Subsampling is achieved by strided convolutions or pooling.

data vector x and a target vector y to train the blackbox system $f(x) \Rightarrow y$. This means that we have to provide extracted time-shifts from other methods, which implicitly introduce assumptions from that method into the supervised model. Alternatively, expensive synthetic models would be required.

The supervised methods are largely based on Optical Flow methods [31], [32]. The FlowNet [31] architecture is based on an Encoder–Decoder CNN architecture. Particularly, FlowNet has reached wide reception and several modifications were implemented, namely FlowNet 2.0 [33] improving accuracy, and LiteFlowNet [34] reducing computational cost. SpyNet [32] and pyramid, warping, and cost volume (PWC)-Net [35] implement stacked coarse-to-fine networks for residual flow correction. PatchBatch [36] and deep discrete flow [37] implement Siamese Networks [38] to estimate optical flow. Alternatively, DeepFlow [39] attempts to extract large displacements optical flow using pyramids of scale-invariant feature transform (SIFT) features. These methods introduce varying types of network architectures, optimizations, and losses that attempt to solve the optical flow problem in computer vision.

2) *Unsupervised CNNs*: Unsupervised or self-supervised CNNs only rely on the data, relaxing the necessity for ground truth time shifts. In [40] the FlowNet architecture is reformulated into an unsupervised optical flow estimator with bidirectional census loss called UnFlow. The UnFlow network relies on the smooth estimation of the forward and backward loss, then adds a consistency loss between the forward and backward loss, and finally warps the monitor to the base image to obtain the final data loss. Optical flow has historically underperformed on seismic data, due to both smoothness and illumination constraints. However, UnFlow replaces the commonly used illumination loss by a ternary census loss [41] with the ϵ -modification by Stein [42]. While this bears possible promise for seismic data, UnFlow implements 2-D losses as opposed to a 3-D implementation that we focus on.

3) *Cycle-Consistent Generative Adversarial Networks*: Cycle-GANs are an unsupervised implementation of generative adversarial networks that are known for domain adaptation [43]. These implement two GAN networks that perform a forward and backward operation that implements a cycle-consistent loss in addition to the GAN loss. The warping problem can be reformulated as a domain adaptation problem. This implements two Generator networks F and G and the

corresponding discriminators D_X and D_Y . These perform a mapping $G : X \rightarrow Y$ and $F : Y \rightarrow X$, trained via the GAN discrimination. The cycle-consistency implements $x \rightarrow G(x) \rightarrow F(G(x)) \approx x$ with the backward cycle-consistency being $y \rightarrow F(y) \rightarrow G(F(y)) \approx y$.

Cycle-GANs such as pix2pix [44] separate image data into a content vector and a texture vector, which could bear promise in the seismic domain, adapting a wavelet vector and an interval vector [3]. However, the confounding of imaging effects, changing underlying geology, changing acquisition, etc., makes the separation nonunique. Moreover, extracting the time shift information and conditioning in the GAN is a very complex problem. The recycle-GAN [45] addresses temporal continuity in videos, this is, however, hard to transfer to seismic data, considering the low number of time-steps in a 4-D seismic survey as opposed to videos. Furthermore, the lack of interpretability of GANs at the point of writing, prohibits GANs from replacing many physics-based approaches, like the extraction of time-shifts.

III. METHOD

The Voxelmorph [15] implements a U-net [46] architecture to obtain a dense warp velocity field and subsequently warps the monitor volume to match the base volume. This minimizes assumptions that have to be satisfied for applying optical flow-based methods. Additionally, the Voxelmorph architecture was specifically developed on medical data. Here, we use an advancement of Voxelmorph that includes a variational layer, which introduced uncertainty to the static velocity estimation, developed in [47]. Medical data often have fewer samples, like seismic data, as opposed to popular video datasets, in which FlowNet and derivative architectures are geared toward the application of popular video datasets. A U-net architecture is particularly suited for segmentation tasks and transformations with smaller than usual amounts of data, considering it was introduced on a small biomedical dataset. The short-cut concatenation between the input and output layers stabilizes training and avoids the vanishing gradient problem. It is particularly suited to stable training in this image matching architecture. In Fig. 2, the U-Net is the leftmost stack of layers, arranged in an hourglass architecture with shortcuts. These feed into a variational layer $\mathcal{N}(\mu, \sigma)$, the variational layer is sampled with the reparametrization trick, due to the sampler not being differentiable [14]. The resulting differential flow is integrated using the VecInt layer, which uses Scaling and Squaring [48]. Subsequently, the data are passed into a spatial transformation layer. This layer transforms the monitor cube according to the warp velocity field obtained from the integrated sampler. The result is used to calculate the data loss between the warped image and the base cube.

More formally, we define two 3-D images $\mathbf{b}, \mathbf{m}$ being the base and monitor seismic, respectively. We try to find a deformation field ϕ parameterized by the latent variable z such that $\phi_z : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. The deformation field itself is defined by this ordinary differential equation (ODE) according to [15]

$$\frac{\partial \phi^{(t)}}{\partial t} = v(\phi^{(t)}) \quad (4)$$

Fig. 2. 2-D representation of modified 3-D Voxelmorph architecture to obtain full-scale warp velocity field. The encoder side of the U-Net architecture consists of four consecutive convolutional (orange) and pooling (red) layers, followed by a convolutional bottleneck layer. The decoder of the U-Net architecture consists of four upsampling (blue) and convolutional layers are connected to the respective same size layers in the encoder. The output is passed to two convolutional layers that are sampled by the reparameterization trick, to provide the static velocity field. The field is integrated via scaling and squaring and passed to the spatial transformer layer (green), which transforms the monitor to optimally match the base image, which is enforced by minimizing the MSE of the images.

where t is time, v is the stationary velocity and the following holds true $\phi^{(0)} = \mathbf{I}$. The integration of v over $t = [0, 1]$ provides $\phi^{(1)}$. This integration represents and implements the one-parameter diffeomorphism in this network architecture. The variational Voxelmorph formulation assumes an approximate posterior probability $q_\psi(z|\mathbf{b}; \mathbf{m})$, with ψ representing the parameterization. This posterior is modeled as a multivariate normal distribution with the covariance $\Sigma_{z|m,b}$ being diagonal

$$q_\psi(z|\mathbf{b}; \mathbf{m}) = \mathcal{N}(z, \mu_{z|m,b}, \Sigma_{z|m,b}) \quad (5)$$

the effects of this assumption are explored in [47].

The approximate posterior probability q_ψ is used to obtain the variational lower bound of the model evidence by minimizing the Kullback–Leibler (KL) divergence with $p(z|\mathbf{b}; \mathbf{m})$ being the intractable posterior probability. Following the full derivation in [47], considering the sampling of $z_k \sim q_\psi(z|\mathbf{b}; \mathbf{m})$ for each image pair $(\mathbf{b}, \mathbf{m})$, we compute $\mathbf{m} \circ \phi_{z_k}$ the warped image we obtain the loss

$$\begin{aligned} \mathcal{L}(\psi; \mathbf{b}, \mathbf{m}) &= -\mathbf{E}_q[\log p(\mathbf{b}|z; \mathbf{m})] \\ &\quad + \text{KL}[q_\psi(z|\mathbf{b}; \mathbf{m}) || p_\psi(z|\mathbf{b}; \mathbf{m})] \\ &\quad + \text{const} \\ &= \frac{1}{2\sigma^2 K} \sum_k \|\mathbf{b} - \mathbf{m} \circ \phi_{z_k}\|^2 \\ &\quad + \frac{1}{2} [\text{tr}(\lambda \mathbf{D} \Sigma_{z|x,y}) - \log \Sigma_{z|x,y}) \\ &\quad + \mu_{z|m,b}^T \Lambda_z \mu_{z|m,b}] + \text{const} \end{aligned} \quad (6)$$

where Λ_z is a precision matrix, enforcing smoothness by the relationship $\Sigma_z^{-1} = \Lambda_z = \lambda \mathbf{L}$, λ controlling the scale of the velocity field. Furthermore, following [47] $\mathbf{L} = \mathbf{D} - \mathbf{A}$ is the Laplacian of a neighborhood graph over the voxel grid,

where $\mathbf{D}$ is the graph degree matrix, and $\mathbf{A}$ defining the voxel neighborhood. K signifies the number of samples. We can express $\mu_{z|m,b}$ and $\Sigma_{z|m,b}$ as variational layers in a neural network and sample from the distributions of these layers. Given the diagonal constraint on Σ , we define the variational layer as the according standard deviation σ of the corresponding dimension. Therefore, we sample $\mathcal{X} \sim \mathcal{N}(\mu, \sigma^2)$ using the reparameterization trick first implemented in variational auto-encoders [49]. The reparameterization trick defines a differentiable estimator for the variational lower bound, replacing the stochastic, nondifferentiable, and therefore, untrainable, sampler.

Defining the architecture and losses as presented in [47], ensures several benefits. The registration of two images is domain-agnostic, which enables us to apply the medical algorithm to seismic data. The warp field is diffeomorphic, which ensures physically viable, topology-preserving warp velocity fields. Moreover, this method implements a variational formulation based on the covariance of the flow field. 3-D warping with uncertainty measure has not been used in seismic data before.

The network is implemented using Keras [50]. Our implementation is based on the original code in the Voxelmorph package [51].

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Experimental Setup

The experimental setup for this article is based on a variation of the modified Voxelmorph [15] formulation. We extended the network to accept patches of data, because our seismic cubes are generally larger than the medical brain scans and therefore,

Fig. 3. Training losses over time with the KL-divergence at the sampling layer, the data loss calculated by MSE, and the combined total loss.

exceed the memory limits of our GPUs. Moreover, Voxelmorph in its original formulation provides subsampled flow fields, this is due to computational constraints. We decided to modify the network to provide full-scale flow fields, despite the computational cost. This enables direct interpretation of the warp field, which is common in 4-D seismic analysis. However, we do provide an analysis in Section IV-B4 of the subsampled flow-field interpolated to full scale, in the way it would be passed to the spatial transformer layer.

The code is made available in [52]. The model is trained with the Adam optimizer [53] with a learning rate of 0.001 and weight decays $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We train the model for 350 epochs to account for experimentation and time. We set the regularization parameter $\lambda = 10$ and the image noise parameter $\sigma = 0.02$ in accordance with the authors of [47]. We adjust the batch size to the maximum on our architecture, which was 16, and purely manually tuned to the maximum possible. The KL divergence and MSE loss are unweighted in the total loss.

The network definition for the subsampled flow field differs from the definition in Fig. 2 that the last upsampling and convolution layer in the Unet, including the skip connection, right before the variational layers (μ, σ) is omitted. That leaves the flow field at a subsampled map by a factor of two. Computationally, this lowers the cost of the Integration operation before resampling for the spatial transformer.

The data situation for this experiment is special in the sense that the method is self-supervised. We, therefore, do not provide a validation dataset during training. The data are 6 surveys from the North Sea. Main field from years 1088, 2005 A, 2005 B, and 2012. Furthermore, we compare to a different field 1903 and 2005 with different geology, acquisition geometry, and acquisition parameters. While we would be content with the method working on the field data (years 1988 and 2005 Survey A) by itself, we do validate the results on separate data from the same field which was acquired with different acquisition parameters and at different times (years 2005 Survey B and 2012). Moreover, we test the data

on seismic data from an adjacent field that was acquired independently (years 1993 and 2005). All data are presented with a relative coordinate system due to confidentiality, where 0 s on the y-axis does not represent the actual onset of the recording. The field geology and therefore, seismic responses are very different. Due to lack of availability, we do not test the trained network on land data or data from different parts of the world. Considering, that the training set is one 4-D seismic monitor-base pair, a more robust network would emerge from training on a variety of different seismic volumes.

Fig. 3 shows the training losses of the batch training. Within a few epochs the network converges strongly, however, within 10 epochs the KL divergence increases slightly over the training. The data loss, optimizing the warping result decreases over the training period. An increase of the KL divergence is acceptable as long as the total loss decreases, which indicates better matching of the volumes. In case the KL divergence would increase vastly, it would violate the base assumption that the static velocity can be approximated by Gaussians and requires reevaluation.

B. Results and Discussion

The network presented generates warp fields in three dimensions as well as uncertainty measures. We present results for three cases in Figs. 4, 9, and 11 with the corresponding warp fields and uncertainties in Figs. 5, 10, and 12. In Fig. 4, we show the results on the data, which the unsupervised method was trained on. Obtaining a warp field on the data itself is a good result, however, we additionally explore the generalizability of the method. Considering the network is trained to find an optimum warp field for the data it was originally trained on, we go on to test the network on data from the same field, that was recorded with significantly different acquisition parameters in Fig. 9. These results test the networks generalizability on colocated data, therefore, not expecting vastly differing seismic responses from the subsurface itself. The imaging differences and differences in equipment in addition to the 4-D difference, however. In Fig. 11, we use the network on unseen data from a different field. The geometry of the field, as well as the acquisition parameters, is different, making generalization a challenge.

In Fig. 4, we collect six 2-D panels from the 3-D warping operation. In Fig. 4(a) and (b), we show the unaltered base and monitor, respectively. The difference between the unaltered cubes is shown in Fig. 4(e). In Fig. 4(c), we show the warped result by applying the z -warp field in Fig. 4(d), as well as the warp fields in (x, y) direction fully displayed in Fig. 5 including their respective uncertainties. The difference of the warped result in Fig. 4(f) is calculated from the matched monitor in Fig. 4(c) and the base in Fig. 4(a).

It is apparent that the matched monitor significantly reduced noise by misaligned reflections. In Table I, we present the numeric results. These were computed on the 3-D cube for an accurate representation. We present the root mean square (rms) and mean absolute error (MAE) and the corresponding difference between monitor and matched difference results. We present rms and MAE to make the values

Fig. 4. Warp results and change in difference on training recall of 1988 to 2005a data. Axes are relative to comply with confidentiality. (a) Base seismic. (b) Monitor seismic. (c) Matched monitor. (d) Z-direction shifts. (e) Difference monitor-base. (f) Difference matched-base.

comparable in magnitude as opposed to the mean squared error (MSE). We present both values, because the rms value is more sensitive to large values, while MAE scales the error linearly, therefore, not masking low amplitude misalignments.

Fig. 5. Warp fields (Top) with uncertainties (Bottom) that accompanies training recall in Fig. 4. (a) X-direction shifts. (b) Y-direction shifts. (c) Z-direction shifts. (d) X-direction uncertainty. (e) Y-direction uncertainty. (f) Z-direction uncertainty.

Both measurements show a reduction in the train data to 50% and the test data on another field show a similar reduction, or below. The test on both the validation data on the same field while the absolute error differs in a stable manner.

TABLE I

QUANTITATIVE EVALUATION OF RESULTS. RMS AND MAE CALCULATED AGAINST RESPECTIVE BASE DATA. TRAINING RECALL, TEST A—SAME FIELD, DIFFERENT ACQUISITION, TEST B—DIFFERENT FIELD, DIFFERENT ACQUISITION

Run	Monitor RMS	Matched RMS	Ratio %	Monitor MAE	Matched MAE	Ratio %
Baseline	0.1047	0.0718	68.6	0.0744	0.0512	68.7
Train	0.1047	0.0525	50.1	0.0744	0.0348	46.7
Test A	0.0381	0.0237	62.2	0.0291	0.0172	59.1
Test B	0.0583	0.0361	62.0	0.0451	0.0254	56.4

In Fig. 5, we present the three-dimensional warp field to accompany the results in Fig. 4. Fig. 5(a)–(c) shows the warp field in x -, y -, and z -direction. The z -direction is generally referred to as time shifts in 4-D seismic. Fig. 5(d)–(f) contains the corresponding uncertainties in x -, y -, and z -direction obtained from the network.

1) *Recall to Training Data:* In Fig. 4, we evaluate the results of the self-supervised method on the training data itself. The main focus is on the main reflector in the center of the panels. The difference in Fig. 4(e) shows that the packet of reflectors marked reservoir in the monitor is out of alignment, causing a large difference, which is corrected for in Fig. 4(f). The topmost section in the panel of Fig. 4(c) shows the alignment of a faulted segment, marked fault in the monitor, to an unfaulted segment in the base. The fault appearing is most likely due to vastly improved acquisition technology for the monitor.

The warp fields in Fig. 5 are an integral part of quality control (QC)-ing the validity of the results. Physically, we expect the strongest changes in the z -direction in Fig. 5(c). The changes in Fig. 5(a) and (b) show mostly subsampling magnitude shifts, except for the x -direction shifts around the fault in the top-most panel present in the monitor in Fig. 4(b). Fig. 5(a) and (b) shows strong shifts at 0.4 s on the left of the panel which corresponds to the strong amplitude changes in the base and monitor. On the one side, these correspond to the strongest difference section, additionally, these are geological hinges, which are under large geomechanical strain. However, these are very close to the sides of the warp, which may cause artifacts. Fig. 5(d)–(f) shows the uncertainty of the network. These uncertainties are across the bank within the 10% range of the sampling rate ($\Delta t = 4$ ms, $\Delta x, y = 25$ m). The certainty within the bulk package in the center of the panels is the lowest in x -, y -, and z -directions. While being relatively lower in the problematic regions discussed before.

The warp field in Fig. 5(d) contains some reflector-shaped warp vectors around 0.4 s, which is due to the wavelet mismatch of the 1988 base to the 2005 monitor. The diffeomorphic nature of the network aligns the reflectors in the image, which causes some reflector artifacts in the z -direction maps.

2) *Comparison to Baseline Method:* We use the dynamic image warping (DIW) method [17] to align the images in Fig. 4. This method extends the DTW method to 2-D and provides a much improved result in 2-D compared to standard cross correlation and DTW methods. Inversion methods need prestack seismic data, which is not available. We chose this baseline to provide a fair comparison with the available data.

Fig. 6. (Left) Comparison of Voxelmorph warp field and (Right) DIW warp fields. (a) Full-scale warp field. (b) DIW timeshifts.

Fig. 6 shows the timeshifts or warp fields generated by the Voxelmorph network and by the DIW algorithm. The DIW algorithm shows a smoothed image. Overall, Fig. 6(b) shows the general trends of Fig. 6(a). The Voxelmorph algorithm is more detailed than the DIW image, however, the general magnitude of the time shifts matches well in the correct areas.

Fig. 7 shows the matched monitors from Voxelmorph and DIW. The matched monitors align quite well without any significant discrepancies. The matched difference shows that the Voxelmorph algorithm performs similar to the baseline method, while removing more 4-D noise from the image. It keeps the 4-D signal intact, albeit slightly varying. The DIW algorithm seems to struggle to align the topmost part of the image, while Voxelmorph aligns these well, removing additional 4-D noise. Table I confirms this quantitatively, where the overall RMSE and MAE are reduced proportionally.

3) *Generalization of the Network:* While the performance of the method on a dataset by itself is good, obtaining a trained model that can be applied on other similar datasets is essential even for self-supervised methods. We test the network on two test sets, Test A is conducted on the same geology with unseen data from a different acquisition, while Test B is on a different field and a different acquisition. The network was trained on a single acquisition relation (2005a–1988). In Fig. 9, we present the crossline data from the same field the network was trained on. The datasets were, however, acquired at different calendar times (2005b–2012), with different acquisition parameters. It follows that although the geology, and therefore, the reflection geometry is similar, the wavelet, and hence, the seismic response is vastly

Fig. 7. Results of Voxelmorph warping compared to baseline DIW algorithm. (Top row) shows the aligned monitor. (Bottom row) shows the difference to the base volume. (a) Base Seismic. (b) DTW monitor seismic. (c) Voxelmorph matched monitor. (d) Unaligned difference. (e) DTW difference. (f) Voxelmorph difference.

different. This becomes apparent when comparing the base Figs. 4(b)–9(a), which were acquired in the same year.

Test A evaluates the network performance on unseen data in the same field (Train: 1988–2005a, Test A: 2005b–2012).

The quantitative results in Table I for Test A generally show lower absolute errors compared to the training results in Section IV-B1. The reduction of the overall amplitudes in the difference maps is reduced by 40%. The unaligned monitor difference in Fig. 9(e) shows a strong coherent difference below the main packet of reflectors around 0.3 to 0.4 s. This would suggest a velocity draw-down in this packet. Although the top half of the unaligned difference contains some misalignment, we would expect the warp field to display a shift around 0.35 s, which can be observed in Fig. 9(d). The aligned difference in Fig. 9(f) contains less coherent differences. The difference does show some overall noise in the maps. This could be improved upon by a more diverse training set. The higher resolution data from 2005 and 2012 possibly have an influence on the result too. Regardless, we can see some persisting amplitude difference around 0.4 s which appears to be a signal as opposed to some misalignment noise above. The warp fields in Fig. 10 show relatively smooth warp fields in x - and y -direction. The warp field in Fig. 10(f) shows overall good coherence, including the change around 0.4 s we would expect. The uncertainty values are in the subsampling range, with the strongest certainty within the strong reflector packet at 0.35 s.

Test B evaluates the network performance on a different field, with different geology, with unrelated acquisition geometry and equipment, and at different times. The test shows a very similar reduction of overall errors in Table I. The rms is reduced by 38% and the MAE is reduced more slightly more in comparison to Test A. In Fig. 11, we present the seismic panels to accompany Test B. The data in Fig. 11(a) and (b) are well resolved and show good coherence. However, the unaligned difference in Fig. 11(e) shows very strong variations in the difference maps. Fig. 11(f) reduces these errors significantly, bringing out coherent differences in the main reflector at 0.27 s. We can see strong chaotic differences in Fig. 11(e) due to the faulted nature of the geology. The network aligns these faulted blocks relatively well, however, some artifacts persist. This is consistent with the warp fields in Fig. 12. The x - and y -directions in Fig. 12(d) and (e), respectively, show overall smooth changes, around faults, these changes are stronger. The z -direction changes are consistent with the training validation and Test A, where the changes are overall stronger. This is also consistent with our geological intuition.

4) *Subsampled Flow*: The original Voxelmorph implementation uses a subsampled warp field. The authors claim two benefits, namely a smoother warp velocity field and reduced computational cost. The aforementioned results were obtained using our full-scale network. In Fig. 8, we present the full scale and upsampled results on the training set. The matched difference in Fig. 8(b) contains more overall noise compared to Fig. 8(a). This is congruent with the warp fields in the figure. The upsampled z -direction warp field in Fig. 8(d) seems to have some aliasing on the diagonal reflector around 0.4 s. This explains some of the artifacts in the difference in Fig. 8(b). The overall warp velocity in Fig. 8(d) is smoother compared to the full-scale field. However, the general structure of coherent negative and positive areas matches in both warp fields, while the details differ. The main persistent difference

Fig. 8. (Top) Comparison of matched differences and (Bottom) z -direction warp field of full-scale neural architecture (Left) and subsampled neural architecture (Right). (a) Full-scale matched difference. (b) Upsampled matched difference. (c) Full-scale warp field. (d) Upsampled warp field.

of the reflector packet at 0.4 s seems similar, nevertheless, the differences further upslope to the right are smoother in the full-scale network result and have stronger residual amplitudes in the upsampled network. Overall, the full-scale network

Fig. 9. Matched difference and warp field for generalization of network to same field with different data (2005b and 2012). (a) Base seismic. (b) Monitor seismic. (c) Matched monitor. (d) Z-direction shifts. (e) Difference monitor-base. (f) Difference matched-base.

results are better for seismic data at a slightly increased computational cost. The subsampled field introduced artifacts in our observations.

We evaluate our network on the training data and two different independent test sets. We do not expect the aligned difference to be exactly zero, due to actual physical changes in

Fig. 10. Warp fields (Top) with uncertainties (Bottom) that accompanies same field generalization in Fig. 9. (a) X-direction shifts. (b) Y-direction shifts. (c) Z-direction shifts. (d) X-direction uncertainty. (e) Y-direction uncertainty. (f) Z-direction uncertainty.

the imaged subsurface. Although the network is unsupervised, a transfer to unseen data is desirable and despite some increase in the overall error possible. The warping on the training data is very good and the warp fields are coherent and reflect

Fig. 11. Matched difference and warp field for generalization of network to a different field (1993 and 2005). (a) Base seismic. (b) Monitor seismic. (c) Matched monitor. (d) Z-direction shifts. (e) Difference monitor-base. (f) Difference matched-base.

the physical reality one would expect. The transfer to unseen data works well, although the misalignment error increases.

Overall, the decrease in both rms and MAE is consistent across test sets.

Fig. 12. Warp fields (Top) with uncertainties (Bottom) that accompanies generalization to different field in Fig. 11. (a) X-direction shifts. (b) Y-direction shifts. (c) Z-direction shifts. (d) X-direction uncertainty. (e) Y-direction uncertainty. (f) Z-direction uncertainty.

V. CONCLUSION

We introduce a deep learning-based self-supervised 4-D seismic warping method. Currently, time shifts are most commonly estimated in 1-D due to computational constraints.

We explore 3-D time-shift estimation as a viable alternative, which decouples imaging and acquisition effects, geomechanical movement, and changes in physical properties like velocity and porosity from confounding into a single dimension.

Existing 3-D methods are computationally expensive, where this learned model can generalize to unseen data without retraining, with calculation times within minutes on consumer hardware. Moreover, this method supplies invertible, reproducible, dense 3-D alignment providing warp fields with uncertainty measures while leveraging recent advancements in neural networks and deep learning.

On the data presented, we obtain subsample scale warp-fields and uncertainties across all directions. We show that the network handles field data with discontinuities and varying seismic signatures well in both training recall and test data. The warp field from this method is less smooth than the conventional method. We show an improvement of full-scale networks over an upsampled network, presented in the original medical Voxelmorph architecture.

ACKNOWLEDGMENT

The authors would like to thank Technical University of Denmark (DTU) Compute for access to the GPU Cluster; and also like to thank Total Exploration and Production Denmark for permission to use the data and publish examples.

REFERENCES

- [1] C. MacBeth, M.-D. Mangriotis, and H. Amini, "Post-stack 4D seismic time-shifts: Interpretation and evaluation," *Geophys. Prospecting*, vol. 67, no. 1, pp. 3–31, 2019.
- [2] J. S. Dramsach, "70 years of machine learning in geoscience in review," in *Machine Learning in Geosciences*. Amsterdam, The Netherlands: Elsevier, 2020, pp. 1–55, doi: [10.1016/bs.agph.2020.08.002](https://doi.org/10.1016/bs.agph.2020.08.002).
- [3] L. Mosser, W. Kimmman, J. Dramsach, S. Purves, A. De la Fuente Briceño, and G. Ganssle, "Rapid seismic domain transfer: Seismic velocity inversion and modeling using deep generative neural networks," in *Proc. 80th EAGE Conf. Exhib. (EAGE)*, Jun. 2018, pp. 1–5.
- [4] S. Li *et al.*, "Deep-learning inversion of seismic data," 2019, *arXiv:1901.07733*. [Online]. Available: <https://arxiv.org/abs/1901.07733>
- [5] B. Liu, S. Yang, Y. Ren, X. Xu, P. Jiang, and Y. Chen, "Deep-learning seismic full-waveform inversion for realistic structural models," *Geophysics*, vol. 86, no. 1, pp. R31–R44, Jan. 2021.
- [6] O. M. Saad and Y. Chen, "Deep denoising autoencoder for seismic random noise attenuation," *Geophysics*, vol. 85, no. 4, pp. V367–V376, Jun. 2020.
- [7] Y. Chen, M. Zhang, M. Bai, and W. Chen, "Improving the signal-to-noise ratio of seismological datasets by unsupervised machine learning," *Seismol. Res. Lett.*, vol. 90, no. 4, pp. 1552–1564, May 2019.
- [8] L. Yang, W. Chen, W. Liu, B. Zha, and L. Zhu, "Random noise attenuation based on residual convolutional neural network in seismic datasets," *IEEE Access*, vol. 8, pp. 30271–30286, 2020.
- [9] J. S. Dramsach, M. Lüthje, and A. N. Christensen, "Complex-valued neural networks for machine learning on non-stationary physical data," *Comput. Geosci.*, vol. 146, Jan. 2021, Art. no. 104643.
- [10] J. S. Dramsach and M. Lüthje, "Deep-learning seismic facies on state-of-the-art CNN architectures," in *Proc. SEG Tech. Program Expanded Abstr.* Tulsa, OK, USA: Society of Exploration Geophysicists, Aug. 2018, pp. 2036–2040.
- [11] G. Zhang, Z. Wang, and Y. Chen, "Deep learning for seismic lithology prediction," *Geophys. J. Int.*, vol. 215, no. 2, pp. 1368–1387, Nov. 2018.
- [12] A. U. Waldeland, A. C. Jensen, L.-J. Gelius, and A. H. S. Solberg, "Convolutional neural networks for automated seismic interpretation," *Lead. Edge*, vol. 37, no. 7, pp. 529–537, Jul. 2018.
- [13] A. Wirgin, "The inverse crime," 2004, *arXiv:math-ph/0401050*. [Online]. Available: <https://arxiv.org/abs/math-ph/0401050>
- [14] D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2575–2583.
- [15] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "VoxelMorph: A learning framework for deformable medical image registration," *IEEE Trans. Med. Imag.*, vol. 38, no. 8, pp. 1788–1800, Aug. 2019.
- [16] P. Hatchell and S. Bourne, "Rocks under strain: Strain-induced time-lapse time shifts are observed for depleting reservoirs," *Lead. Edge*, vol. 24, no. 12, pp. 1222–1225, Dec. 2005.
- [17] D. Hale, "Dynamic warping of seismic images," *Geophysics*, vol. 78, no. 2, pp. S105–S115, Mar. 2013.
- [18] S. Fomel and L. Jin, "Time-lapse image registration using the local similarity attribute," *Geophysics*, vol. 74, no. 2, pp. A7–A11, 2009.
- [19] E. Zabihi Naeini, H. Hoerber, G. Poole, and H. R. Siahkoobi, "Simultaneous multivintage time-shift estimation," *Geophysics*, vol. 74, no. 5, pp. V109–V121, Sep. 2009.
- [20] C. MacBeth, H. Amini, and S. Izadian, "Review paper: Methods of measurement for 4D seismic post-stack time shifts," *Geophys. Prospecting*, vol. 68, no. 9, pp. 2637–2664, Sep. 2020.
- [21] S. A. Hall, C. MacBeth, O. I. Barkved, and P. Wild, "Cross-matching with interpreted warping of 3D streamer and 3D ocean-bottom-cable data at Valhall for time-lapse assessment," *Geophys. Prospecting*, vol. 53, no. 2, pp. 283–297, Mar. 2005.
- [22] D. Hale, "An efficient method for computing local cross-correlations of multi-dimensional signals," *CWP Rep.*, 2006, p. 656.
- [23] J. Rickett, L. Duranti, T. Hudson, B. Regel, and N. Hodgson, "4D time strain and the seismic signature of geomechanical compaction at genesis," *Lead. Edge*, vol. 26, no. 5, pp. 644–647, May 2007.
- [24] A. J. Cherrett, I. Escobar, and H. P. Hansen, "Fast deterministic geostatistical inversion," in *Proc. 73rd EAGE Conf. Exhib. Incorporating SPE EUROPEC*, May 2011, p. 238.
- [25] B. Zitová and J. Flusser, "Image registration methods: A survey," *Image Vis. Comput.*, vol. 21, no. 11, pp. 977–1000, Oct. 2003.
- [26] A. Goshtasby, "Image registration by local approximation methods," *Image Vis. Comput.*, vol. 6, no. 4, pp. 255–261, Nov. 1988.
- [27] R. Bajcsy and S. Kovačič, "Multiresolution elastic matching," *Comput. Vis., Graph., Image Process.*, vol. 46, no. 1, pp. 1–21, Apr. 1989.
- [28] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. Pluim, "Elastix: A toolbox for intensity-based medical image registration," *IEEE Trans. Med. Imag.*, vol. 29, no. 1, pp. 196–205, Jan. 2010.
- [29] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "3D brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.*, vol. 39, no. 3, p. 609, 1994.
- [30] M. F. Beg, M. I. Miller, A. Trounev, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. Comput. Vis.*, vol. 61, no. 2, pp. 139–157, Feb. 2005.
- [31] A. Dosovitskiy *et al.*, "FlowNet: Learning optical flow with convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2015, pp. 2758–2766.
- [32] A. Ranzan and M. J. Black, "Optical flow estimation using a spatial pyramid network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 4161–4170.
- [33] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 2462–2470.
- [34] T.-W. Hui, X. Tang, and C. C. Loy, "LiteFlowNet: A lightweight convolutional neural network for optical flow estimation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8981–8989.
- [35] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8934–8943.
- [36] D. Gadot and L. Wolf, "PatchBatch: A batch augmented loss for optical flow," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 4236–4245.
- [37] F. Güney and A. Geiger, "Deep discrete flow," in *Proc. Asian Conf. Comput. Vis. Cham, Switzerland: Springer*, 2016, pp. 207–224.
- [38] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. CVPR*, vol. 1, 2005, pp. 539–546.
- [39] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid, "DeepFlow: Large displacement optical flow with deep matching," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 1385–1392.
- [40] S. Meister, J. Hur, and S. Roth, "Unflow: Unsupervised learning of optical flow with a bidirectional census loss," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018, pp. 1–9.
- [41] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," in *Proc. Eur. Conf. Comput. Vis.* Berlin, Germany: Springer, 1994, pp. 151–158.
- [42] F. Stein, "Efficient computation of optical flow using the census transform," in *Joint Pattern Recognit. Symp.* Berlin, Germany: Springer, 2004, pp. 79–86.

- [43] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 2223–2232.
- [44] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 1125–1134.
- [45] A. Bansal, S. Ma, D. Ramanan, and Y. Sheikh, "Recycle-GAN: Unsupervised video retargeting," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 119–135.
- [46] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2015, pp. 234–241.
- [47] A. V. Dalca, G. Balakrishnan, J. Guttag, and M. R. Sabuncu, "Unsupervised learning for fast probabilistic diffeomorphic registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2018, pp. 729–738.
- [48] N. J. Higham, "The scaling and squaring method for the matrix exponential revisited," *SIAM J. Matrix Anal. Appl.*, vol. 26, no. 4, pp. 1179–1193, Jan. 2005.
- [49] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [50] F. Chollet *et al.* *Keras*. (2015). [Online]. Available: <https://keras.io>
- [51] A. V. Dalca *et al.* *Voxelmorph*. (2018). [Online]. Available: <https://voxelmorph.mit.edu>
- [52] J. S. Dramsch, "4D seismic warping voxelmorph code," Figshare Softw., doi: [10.6084/m9.figshare.12808910.v1](https://doi.org/10.6084/m9.figshare.12808910.v1).
- [53] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>

Jesper Sören Dramsch received the B.Sc. degree in geophysics and oceanography and the M.Sc. degree in geophysics from the University of Hamburg, Hamburg, Germany, in 2010 and 2014, respectively. He is pursuing the Ph.D. degree in 4-D seismic and machine learning with the Technical University of Denmark (DTU), Lyngby, Denmark.

He visited Heriot-Watt University, Edinburgh, U.K., as a Visiting Scholar from 2018 to 2019 and taught programming and machine learning principles on a consultant basis in 2019. His research

interests lie in digital signal and image processing, geophysics, as well as machine learning and neural networks, and particularly in geoscience applications.

Anders Nymark Christensen received the B.Sc. and M.Sc. degrees in biomedical engineering from the Technical University of Denmark (DTU), Lyngby, Denmark, in 2009 and 2012, respectively, and the Ph.D. degree in data analysis of medical images from the Department of Applied mathematics and Computer Science, DTU in 2012.

He has worked with image analysis and machine learning in both medical and industrial applications as well as remote sensing. Since 2016, he has been an Assistant Professor with the Section for

Image Analysis and Computer Graphics, DTU. His research interests include applied image analysis, statistics, and machine learning and especially their applications in data rich and label sparse situations.

Colin MacBeth received the B.A. and M.A. degrees in physics from the University of Oxford, Oxford, U.K., in 1980, and the Ph.D. degree in geophysics from the University of Edinburgh, Edinburgh, U.K., in 1983.

He is a Professor of reservoir geophysics with the Institute of Geoenergy Engineering, Heriot-Watt University, Edinburgh, U.K. He has authored or coauthored more than 230 journal articles and conference paper publications in the general areas of seismic reservoir characterization. His main interests

are in quantitative interpretation of 4-D seismic, data integration and seismic history matching.

Dr. MacBeth is a member of the European Association of Exploration Geophysics (EAGE), the Society of Exploration Geophysicists (SEG), the Society of Petroleum Engineers (SPE), the Petroleum Exploration Society of Great Britain (PESGB), and a fellow of the Royal Scottish Society of Arts (SSRA). He was a recipient of the Conrad Schlumberger Award from the EAGE for his contributions to Geoscience and Engineering in 2007 and the Norman Falcon Award for the best paper in Petroleum Geoscience in 2018.

Mikael Luthje received the B.Sc. degree in physics from the Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark, in 2000, and the M.Sc. degree in geophysics from University Centre in Svalbard, Longyearbyen, Svalbard, in 2001, and the Ph.D. degree in applied mathematics and remote sensing from the Department of Electromagnetic Systems (now National Space Institute), Technical University of Denmark, Lyngby, Denmark; the Centre for Polar Observation and Modeling, University College London, London, U.K.; and the Department

of Applied Mathematics and Theoretical Physics, Cambridge University, Cambridge, U.K., in 2005.

He has worked within the industry and applied research in U.K., Norway, and Denmark. He has been a Senior Researcher with the Technical University of Denmark, Lyngby, Denmark, since 2015. His research interests include prediction of fractures and faults for improved numerical flow modeling with a recent focus on CO₂ sequestration and geothermal systems as well as remote sensing and deep learning for geophysical applications.