

Deep Learning Tomography

Mauricio Araya-Polo¹, Joseph Jennings^{1,2}, Amir Adler³, Taylor Dahlke²

Abstract

Velocity model building (VMB) is a key step in hydrocarbon exploration; The VMB main product is an initial model of the subsurface that is subsequently used in seismic imaging and interpretation workflows. Reflection or refraction Tomography and full waveform inversion (FWI) are the most commonly used techniques in VMB. On one hand, Tomography is a time-consuming activity that relies on successive updates of highly human-curated analysis of gathers. On the other hand, FWI is very computationally demanding with no guarantees of global convergence.

We propose and implement a novel concept that bypasses these demanding steps, directly producing an accurate gridding or layered velocity model from shot gathers. Our approach relies on training deep neural networks; the resulting predictive model maps relationships between the data space and the final output (particularly, the presence of high velocity segments that might indicate salt formations). In term of time, the training task takes a few hours for 2D data, but the inference step (predicting a model from previously unseen data) takes only seconds.

The promising results shown here for synthetic 2D data demonstrate a new way of using seismic data and suggests fast turnaround of workflows that now make use of machine learning approaches to identify key structures in the subsurface.

Introduction

Exploration workflows are under great pressure, from improving performance at lower costs to the ongoing avalanche of data coming from new generations of sensors and modern acquisition systems. Some of the key steps in these workflows depend on domain experts, their time is precious and limited, but the amount of data that needs to be thoroughly analyzed is increasing. In addition, the complexity of some of the exploration areas requires extra attention. The problem can be summarized as an explosion of data, which are increasingly more complex than even before.

Geoscientists need to be empowered with new tools, tools that digest, as much of the data is possible before the human expert intervenes. The High Performance Computing (HPC) revolution of the ten years (BizTech, 2014) shares the same purpose but essentially targets processing speed rather than any other specific step of the exploration workflow.

¹ Shell International Exploration & Production Inc.

² Stanford University

³ MIT, Center for Brains, Minds and Machines, e-mail: adleram@mit.edu

This work was supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF - 1231216.

This time around advanced data-oriented algorithms look to improve every step of the workflow through a deeper understanding of the data, from extracting the relevant information to a better awareness of the rest of the steps in a more integrated fashion rather silos of knowledge.

What we propose in this work goes beyond what is becoming the new norm, which is Machine Learning (ML) techniques being applied to specific well-known steps of the workflow. This same magazine carried a special edition on how Analytics and ML (March 2017) techniques are paving inroads in different aspects of the exploration workflow, but still most work is focused on identifying features or attributes in migrated images (Hale, 2012; Hale, 2013; Guillen, 2015; Addison, 2016; Bougher et al, 2016), therefore helping to tackle the interpretation step, very few has been proposed on helping directly processing or VMB. In general, the literature is abundant with refinements to this workflow, but still it remains largely untouched.

Our method produces velocity models directly from raw seismic data in a way that is alternative to classic Tomography, also it is automatic and without human intervention. The ML technique employed follows recent work (Frogner et al, 2015; Zhang et al, 2014; Dahlke et al, 2016, Araya-Polo et al, 2017) that demonstrates this new approach, which uses a Deep Neural Network (DNN) statistical model to transform raw input seismic data directly to the final mapping in 2D or 3D. The computational costs come mostly from training and this happens only once up front. After training, velocity model reconstruction costs are negligible, thus making the overall computing costs a fraction of the time needed for traditional techniques, in particular the ones involving partial differential equations-based simulation. One key element of our method is the use of a feature based on semblance, therefore this feature pre-digest velocity information for the training process, still this feature extraction step is automated and not subjected to human bias.

In terms of deployment modes, we foresee models being trained with specific data belonging to different major formations, such as unconventional, pre-salt or subsalt. The main concerns relate to the generalization error, which basically set the limits on how much a predicting model can accurately predict for unseen data. Finally, regarding exploration workflows, one can imagine this technique being used just after data acquisition, then trained models are loaded up to the cloud from which interpreters can pull realizations, thus performing online scenarios testing when feeding back their model modifications to applications such as the one proposed in Araya-Polo et al, 2017. This imagined workflow is fully ML-based, flexible and with the domain experts at the center of the critical decision making process; if it is accompanied by the proper resourcing this workflow approaches significantly to a real-time ubiquitous experience.

The layout of the paper is as follows; we start by explaining the basics of the problem followed by discussions on Deep Learning. Next, we introduce the general workflow used by our ML system, that we termed GeoDNN. After that, we discuss

our results and experiments with 2D synthetic data. Finally, conclusions are presented which include the directions for future work.

Problem formulation

Formally, the traditional tomography problem can be expressed as the minimization of the following objective function:

$$J(m) = ||d_m(m) - d_{obs}||_2^2,$$

where m is the optimal velocity model that minimizes $J(m)$, d_m is a data vector that is modeled from a non-linear modeling operator $f(m)$, and d_{obs} is the recorded data vector. While it is common to minimize the sum of the squares (represented here by the square of the L2 norm), other objective functions may be used. Note that in the case of travel-time tomography the data vectors contain travel-times that are modeled via the solution of the Eikonal equation. Alternatively, in the FWI case, the data contain the seismic traces that are modeled via the numerical solution of the wave-equation.

As is apparent by the non-linear relationship between d_m and m this inversion is nonlinear. Additionally, since for reflection seismic surveys d_{obs} contains surface seismic data, it does not contain all of the necessary information to define a velocity model that varies arbitrarily with depth and along the horizontal directions (Biondi, 2001). This means that, in general, minimizing the above equation is an ill-posed problem. While in using a deep learning approach to tomography we do not rely on numerical solutions of the Eikonal or wave-equations, we still need to consider the nonlinearity and ill-posed nature of this inverse problem.

The application of neural networks for velocity estimation and for geophysical applications in general is not new (van der Baan, 2000). The first use of neural networks for velocity estimation was proposed by Roth and Tarantola (1994) where neural networks are used to estimate 1D velocity functions from shot gathers. Nath et. al., 1999 used neural networks for travel-time cross-well tomography. After training their network using travel-time maps and synthetic velocity models as training data, the network was then used to tomographically estimate velocities for cross-well data acquired in West Bengal, India. Although the problem we attempt to solve is similar, our work is novel in that it makes use of the recent development of more advanced deep neural network (DNN) architectures and moreover, we use all of the data (not only selected travel-times) in order to train our DNN and perform tomography.

Machine Learning of Tomography Operators via Deep Neural Networks

Using machine learning algorithms is an appealing alternative to classic seismic processing, and among this class of algorithms we have implemented the tomography operator using a Deep Neural Network (DNN). The tomography operator is learned from seismic training data, using statistical learning (Hastie et al, 2001) principles. The tomography process is depicted in Figure 1, and it performs

reconstruction of the velocity model from raw seismic traces, or from features computed from raw seismic traces (as part of the tomography operator). In a real-life application, the ground-truth model is unavailable, and the tomography operator is designed to minimize the difference between the reconstructed velocity model and the (unavailable) ground-truth one.

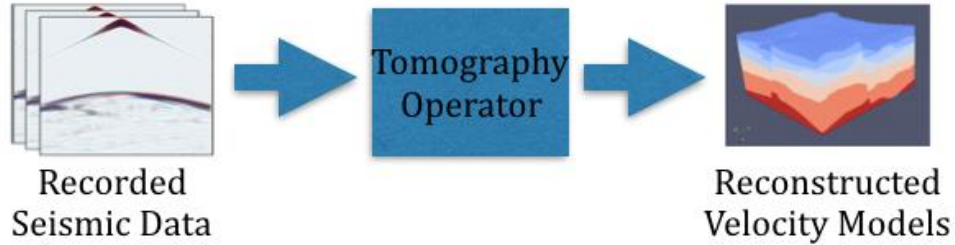


Figure 1: Tomography reconstruction of velocity models from seismic data.

In the statistical learning framework, the tomography operator is learned using a collection of N training examples $\{X_i, V_i\}_{i=1}^N$, where X_i denotes the seismic traces (or features of seismic traces) generated from the i -th velocity model V_i . Specifically, the tomography operator is learned by solving the following optimization problem:

$$\hat{\alpha} = \arg \min_{\alpha} \frac{1}{N} \sum_{i=1}^N L \left(V_i, \underbrace{T(X_i, \alpha)}_{\hat{V}_i} \right),$$

where $T(X_i, \alpha)$ is the tomography operator, parameterized by the coefficients vector α , and its output is the reconstructed velocity model \hat{V}_i . The loss function $L(V_i, \hat{V}_i)$ measures the difference between the ground truth velocity model V_i and its reconstructed version \hat{V}_i . The loss function we employed is the squared error⁴ $L(V_i, \hat{V}_i) = (V_i - \hat{V}_i)^2$, which is frequently used in regression problems, and leads to the following optimization problem:

$$\hat{\alpha} = \arg \min_{\alpha} \frac{1}{N} \sum_{i=1}^N (V_i - T(X_i, \alpha))^2.$$

A frequently used minimization approach is the Gradient Descent (GD) which iteratively updates the coefficients vector as follows:

$$\alpha_{t+1} = \alpha_t - \mu \frac{\partial L_E}{\partial \alpha}$$

where μ is a positive learning rate, L_E is the empirical loss:

⁴Note that in the case of two images, the squared error loss is computed pixel-based, namely, it is the sum of all squared pixels differences.

$$L_E = \frac{1}{N} \sum_{i=1}^N L(V_i, T(X_i, \alpha)),$$

and the gradient of L_E , with respect to α , is given by:

$$\frac{\partial L_E}{\partial \alpha} = \frac{1}{N} \sum_{i=1}^N \frac{\partial L}{\partial \alpha}(V_i, T(X_i, \alpha)).$$

The tomography operator $T(X_i, \alpha)$ was implemented using a Deep Neural Network (DNN), as detailed in the following.

Deep Neural Networks

DNNs are powerful machine learning algorithms (LeCun et al, 2015, Goodfellow et al, 2016), which provide state-of-the-art results in numerous computer vision, speech processing, and artificial intelligence problems. In particular, DNNs provide excellent results for imaging inverse-problems such as de-noising (Burger et al 2012, Xie et al, 2012), super-resolution (Dong et al, 2016), compressed-sensing (Adler et al, 2017), and X-ray computed tomography (Wang, 2016, Würfl et al, 2016). In addition, according to the Universal Approximation Theorem (Hornik et al, 1989), DNNs can be used to approximate any arbitrary continuous function up to a specified accuracy. For these reasons, there is great promise in using this approach to approximate complex functions that are highly non-linear.

DNNs are comprised of “layers” of weighted nodes, as depicted in Figure 2: the input to the network is connected to the input layer, which is followed by a varying number of hidden layers, and eventually the output of the network is computed at the output layer. Each hidden layer’s inputs are activated by the outputs of the previous layer. These networks are trained with examples, per the statistical learning approach where the correct output (label) is known for a given input, and the weight parameters in the nodes of the network evolves due to the minimization of the error between the prediction and true value. This causes the network to increasingly become a better predictor of the training examples, and ultimately of any example (assuming proper training) of a class of data that is similar in nature to the training data.

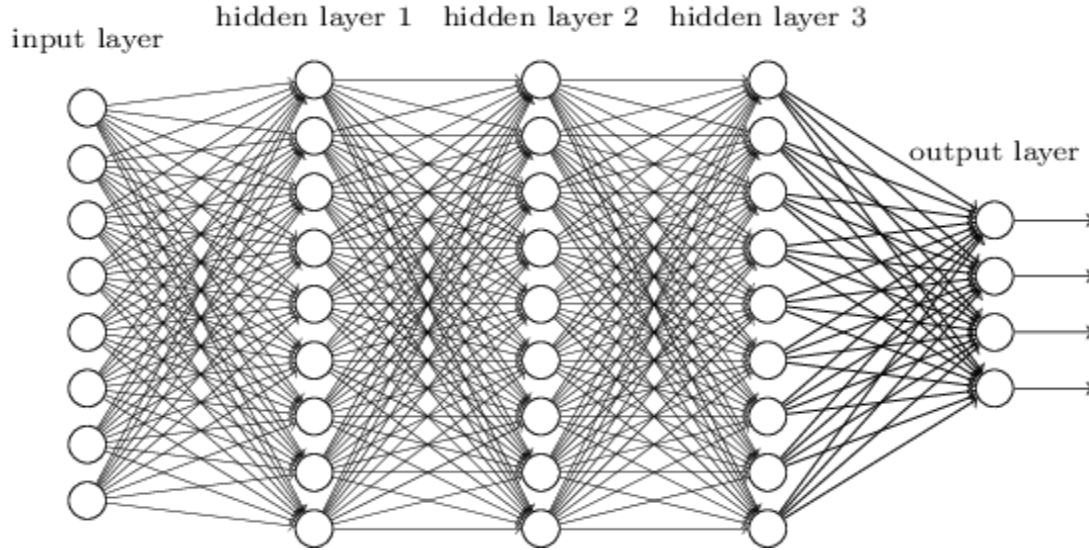


Figure 2: Topology of a Deep Neural Network with three hidden layers.

The proposed tomography operator is therefore described as follows, assuming for example three hidden layers:

$$T(X, \alpha) = f_{out}(f_3(f_2(f_1(X, \alpha_1), \alpha_2), \alpha_3), \alpha_{out}),$$

where f_{out} is the output layer function, parameterized by α_{out} and the hidden layer functions are f_1, f_2, f_3 each parameterized by $\alpha_1, \alpha_2, \alpha_3$, respectively (the vector α is composed by $\alpha_1, \alpha_2, \alpha_3, \alpha_{out}$).

Our ability to design effective neural networks is limited by computing resources constraints. More complex networks are more computational demanding to train, and the generation of accurate training examples can be computationally expensive due to large-scale forward modeling. Ultimately, our predictions are only as good as the complexity and refinement of our neural network coupled with the relevance and quality of the features that we choose as inputs.

Workflow

Since we lacked abundant labeled data, this implies that the neural network's best result would be bound by the limited number of examples, which often leads to over-fitting of the learning procedure to the training data. and control of the main parameters involved (data generation) is key when proving a new concept, in this work we focus on results for 2D synthetic only. We introduce two workflows, one for training and one for inference (aka testing) as explained below.

In the training workflow (Fig 3), the first step is the pseudo-random generation of thousands of unbiased velocity models and from them the labels that represent the experiment, for instance models with faults or salt bodies. In the second step, a modeling step produces the seismic data, for the sake of simplicity and brevity only acoustic approximation of the wave equation is used. The third step extracts

features from the seismic data. The purpose of this step is twofold, on one hand, to reduce the amount of data used for training, which therefore alleviate the stress on the computing resources, on the other hand, it helps the training to focus on certain aspects within the data that are relevant for the experiment, this also helps with the accuracy and convergence of the training task. Once we have extracted the features, the actual deep learning process starts. Our workflow is fully parametrical, from the velocity generation to the feature extraction; therefore, the richness of the experiments is comprehensive in terms of variety of velocity models, acquisition geometries, etc.

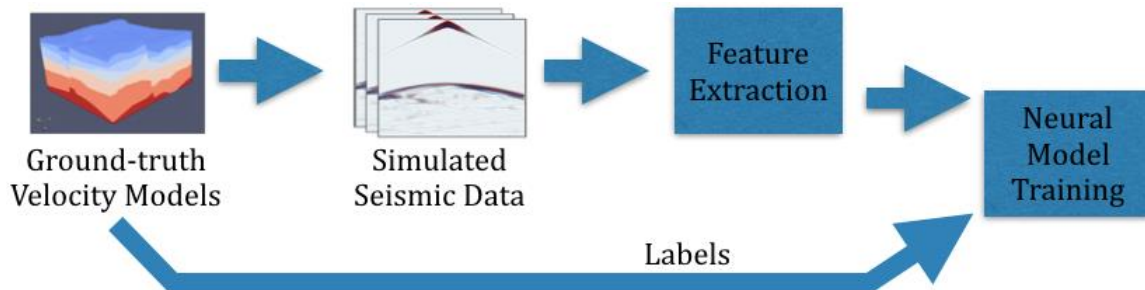


Figure 3. Training workflow.

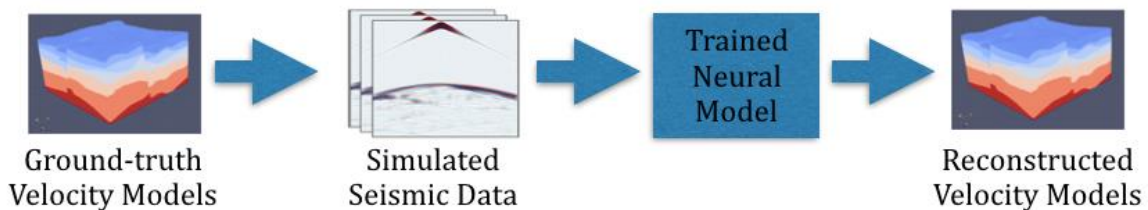


Figure 4. Inference (testing) workflow.

The inference workflow (Figure 4), is the workflow where new models are predicted when exposed to unseen or new data. In our particular context of using synthetic data, it starts in the same fashion as the training workflow, models and data is generated, then that data -that has not been used for training- is presented to the predicting model that reconstruct a velocity model. Since we generate the testing data following the mentioned steps, calculations of accuracy of the model are straightforward.

Semblance as a feature for Machine Learning

Feature extraction is a key step in our workflow as it can greatly improve the training of the DNN by providing it with the most relevant data for learning. Our Machine Learning platform GeoDNN is capable of handling diverse network architectures and data, but given that we desire to learn a velocity tomography operator from the data, we perform velocity analysis and provide semblance panels for different common midpoint (CMP) locations as the input feature. To calculate the semblance panel for a given midpoint, we first apply a normal moveout (NMO)

correction to a common midpoint gather using the second-order travel time equation:

$$t_{NMO}^2 = t_0^2 + \frac{x^2}{V_{NMO}^2}$$

where t_{nmo} is the calculated NMO travel-time, t_0 is the zero-offset travel time, x is the offset and V_{NMO} is the NMO velocity. By choosing a trial V_{NMO} we can then perform an NMO correction on the gather resulting in an NMO-corrected image $q[j, k]$ where j and k are the corrected NMO time and offset sample indices respectively (following the notation of Luo and Hale 2012). Semblance is then calculated by stacking along the offset index and smoothing along the time index of $q[j, k]$. This can be expressed mathematically as:

$$s[i] = \frac{\sum_{j=i-M}^{i+M} (\sum_{k=0}^{N-1} q[j, k])^2}{N \sum_{j=i-M}^{i+M} \sum_{k=0}^{N-1} q[j, k]^2}$$

where $s[i]$ is the output semblance at the output time sample i , N is the total number of offset samples and M is a parameter that defines the length of the time-smoothing window of length $2M + 1$ centered at i . Additionally, we calculate weighting functions that are applied to semblance panels that emphasize terms in the semblance calculation that are most sensitive to changes in velocity (Luo and Hale 2012). While in the semblance calculation we assume for now only second-order moveout (i.e., the traditional NMO equation), we have the capability of using higher-order terms in the travel-time equation allowing for greater accuracy at far offsets (Yilmaz, 2001).

Given that we provide semblance panels for multiple CMP locations, this input feature ends up having three dimensions, making a cube. Figures 5 & 6 show us two things about this feature space. First, for the particular model that the semblance cube represents, we have a high percentage of zero-entry and low value parameters. This is true for many models that we perform semblance cubes on, which means there is an opportunity to sparsify the parameter space. Second, the events in the semblance cube space have patterns that relate to the reflector position and velocity. Humans can interpret some of these patterns (such as distinct energy spikes / clusters which correspond to sharp unpolluted reflection events). Other patterns that are mixed or smeared across the semblance space can imply non-uniqueness which is much more difficult to derive a model approximation from. The advantage of using machine learning is that we are able to leverage the ability of the DNN to learn from a multitude of examples to discover complex patterns that would otherwise be very expensive and difficult to learn and utilize. Using these patterns, we can learn a mapping from the semblance space to the velocity model space. Alternative methods like inversion can be very expensive because the mapping between these spaces uses wave propagation (or some other approximation) as the forward operator. Other methods would try to linearize the forward operator, or follow some more simplistic methodology like picking velocities from peak

amplitudes in the semblance cube. All of these methods need to be repeated for each model of interest. Using a DNN methodology, we need to only train once, after which subsequent model approximations can be found from their corresponding semblance plots at negligible cost.

As stated, the nature of using semblance as a feature input is that there are patterns that have meaning in relation to the velocity model, some of which are trivial to explain while others are more complex. DNN architectures fundamentally learn patterns in the feature space using stencils whose dimensions are predetermined. We believe that the geophysics based transformation that the semblance cube represents make it a good choice as an input feature for deep learning for tomographic velocity estimation, especially for DNN architectures that can leverage the patterns that are found in that space.

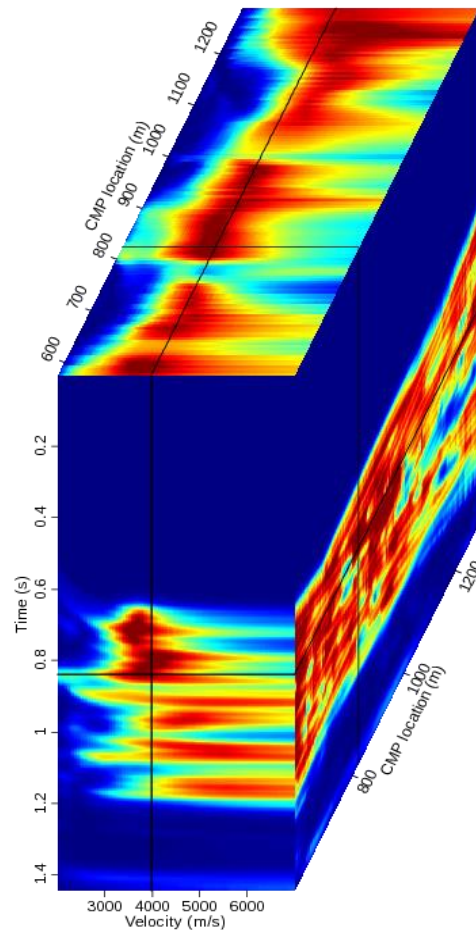


Figure 5. A calculated semblance cube used as an input feature for deep learning. The front face of the cube (with axes of zero-offset time and velocity) shows the semblance panel for a particular CMP location used in traditional velocity analysis. The side face of the cube (with axes of CMP location and zero-offset time) shows the calculated semblance for a particular velocity for all CMP locations and time. Note the spatially coherent structure of the semblance in the cube.

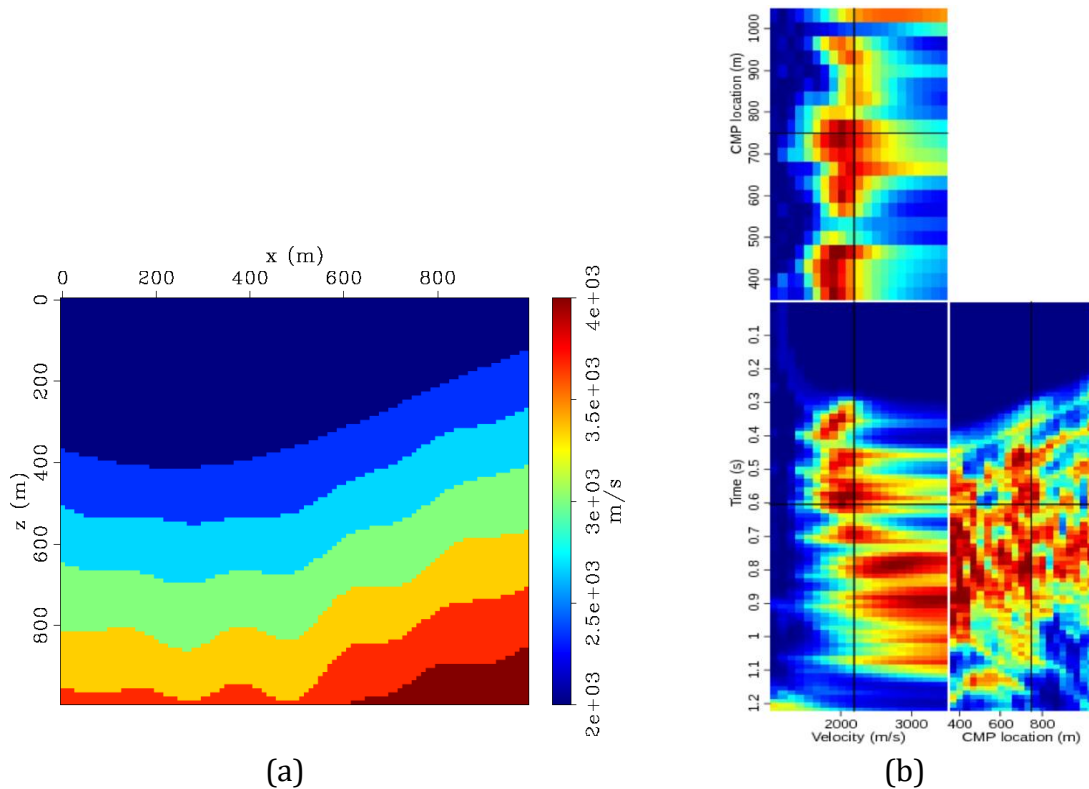


Figure 6. An example of a calculated semblance cube for a seven-layer model. Note that the traditional velocity analysis panel of the semblance cube shown in (b) captures each of the layer interfaces. Also, note that the right panel (axes of CMP location and zero-offset time) of the semblance cube qualitatively gives the approximate structure of the velocity model shown in (a).

Implementation and Results

We generated thousands of random two-dimensional velocity models with up to four faults in them, of varying strike, dip angle, and position. Our models had between three and eight layers each, with velocities varying from 2000 to 4000 [m/s], with layer velocity increasing with depth. These models were 140x180 grid points at the sampling used for wave equation solving. The raw data collected was reduced to a semblance feature set that can fit in multiple NVIDIA K80 GPGPU memory.

We trained the proposed DNN using a training set composed of tens of thousands of velocity models, and tested the tomography results with a testing set of thousands of velocity models.

Experiments type I:

The output of the DNN is continuous valued image, whereas the ground-truth velocity model images are comprised from a discrete number of values, each correspond to a unique velocity value. Therefore, we have applied a post-processing

image segmentation (Szeliski, 2010) stage to each reconstructed velocity model using two methods: (1) K-Means segmentation which uses the ground-truth number of layers in order to cluster all pixels into to correct number of segments. (2) K-Means segmentation with 8 segments (layers) for all velocity models (in a real application, the number of segments can be controlled). The visual quality of each segmented image was compared against the ground truth velocity model (i.e. test example label), using the Structural Similarity Image Metric (SSIM), developed by Wang et. al. (Wang, 2004). The SSIM metric is computed using three image features that mostly influence the human visual system: structure, contrast and luminance. Given two images, the SSIM formula computes a continuous number between 0.0 to 1.0, where 1.0 corresponds to identical images, and 0.0 corresponds to complete visually dissimilar images. The SSIM metric is considered more coherent to human judgment, than the Mean Squared Error (MSE) metric, for image comparisons. The averaged SSIM over thousands of test velocity models is 0.8717 for K-Means with the correct number of segments, and 0.8603 for K-Means with 8 segments for all images, which clearly indicate very high similarity to the ground truth velocity models. In Figures 7, we provide examples of the reconstructed velocity models with varying numbers of layers, which demonstrate the high visual quality of the reconstructed images. We have observed that the reconstruction network tends to smooth faults (third row of Figure 7), and further improvements for accurate faults reconstruction are left for future research.

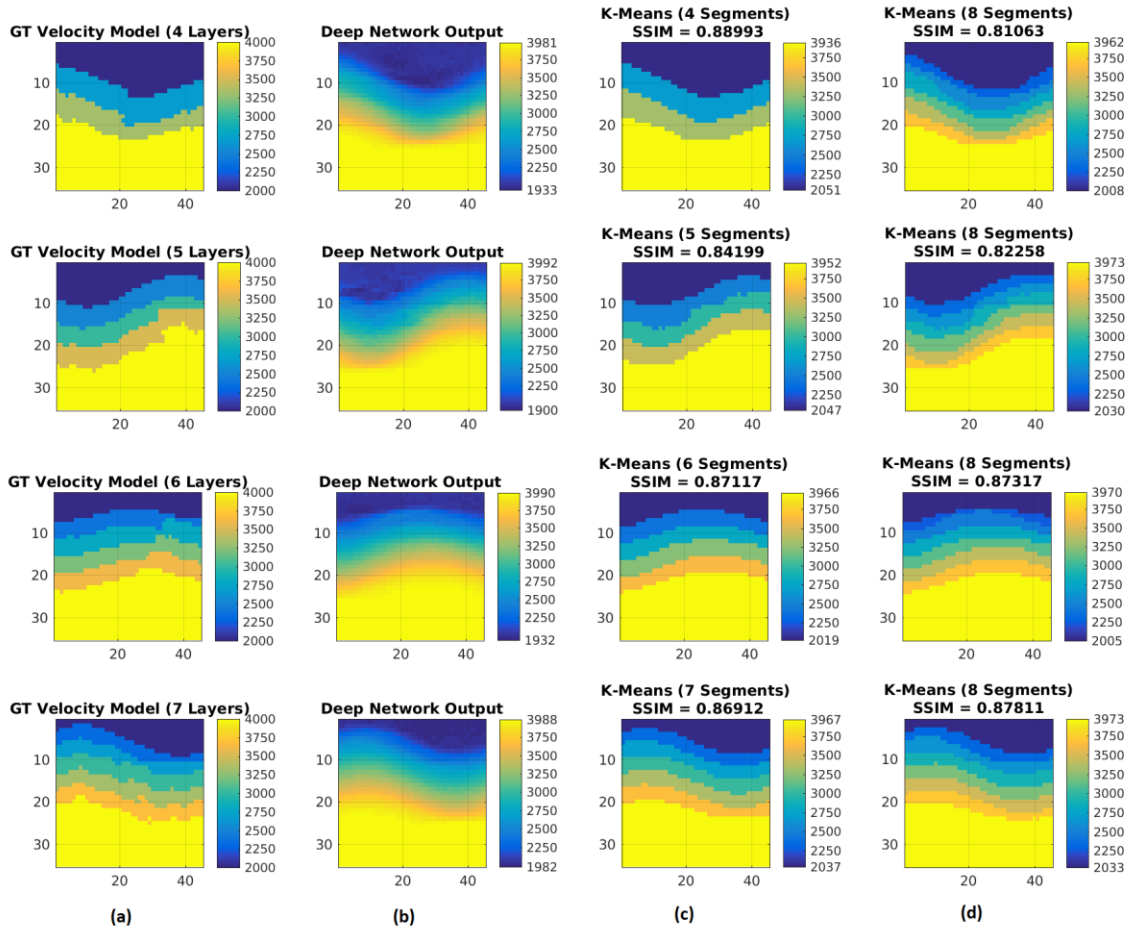


Figure 7. Tomography with 4-7 velocities: (column a) Ground truth (column b) DNN output image (column c) Segmented image using K-Means and the correct number of layers; and (column d) Segmented image using K-Means and 8 layers.

Experiments type II:

In this set of the experiments the labels and reconstructed models are continuous value (not a binary or multiclass classification process) that represents velocity. Some of the labels and models include salt bodies along with 3 to 7 layers. The evaluation metrics are R^2 score (coefficient of determination) and SSIM as described above. R^2 score measures the total variation of the outcomes provided by the model, it is interpreted as the goodness of the model fitting, the values can be negative and the optimal value is 1.

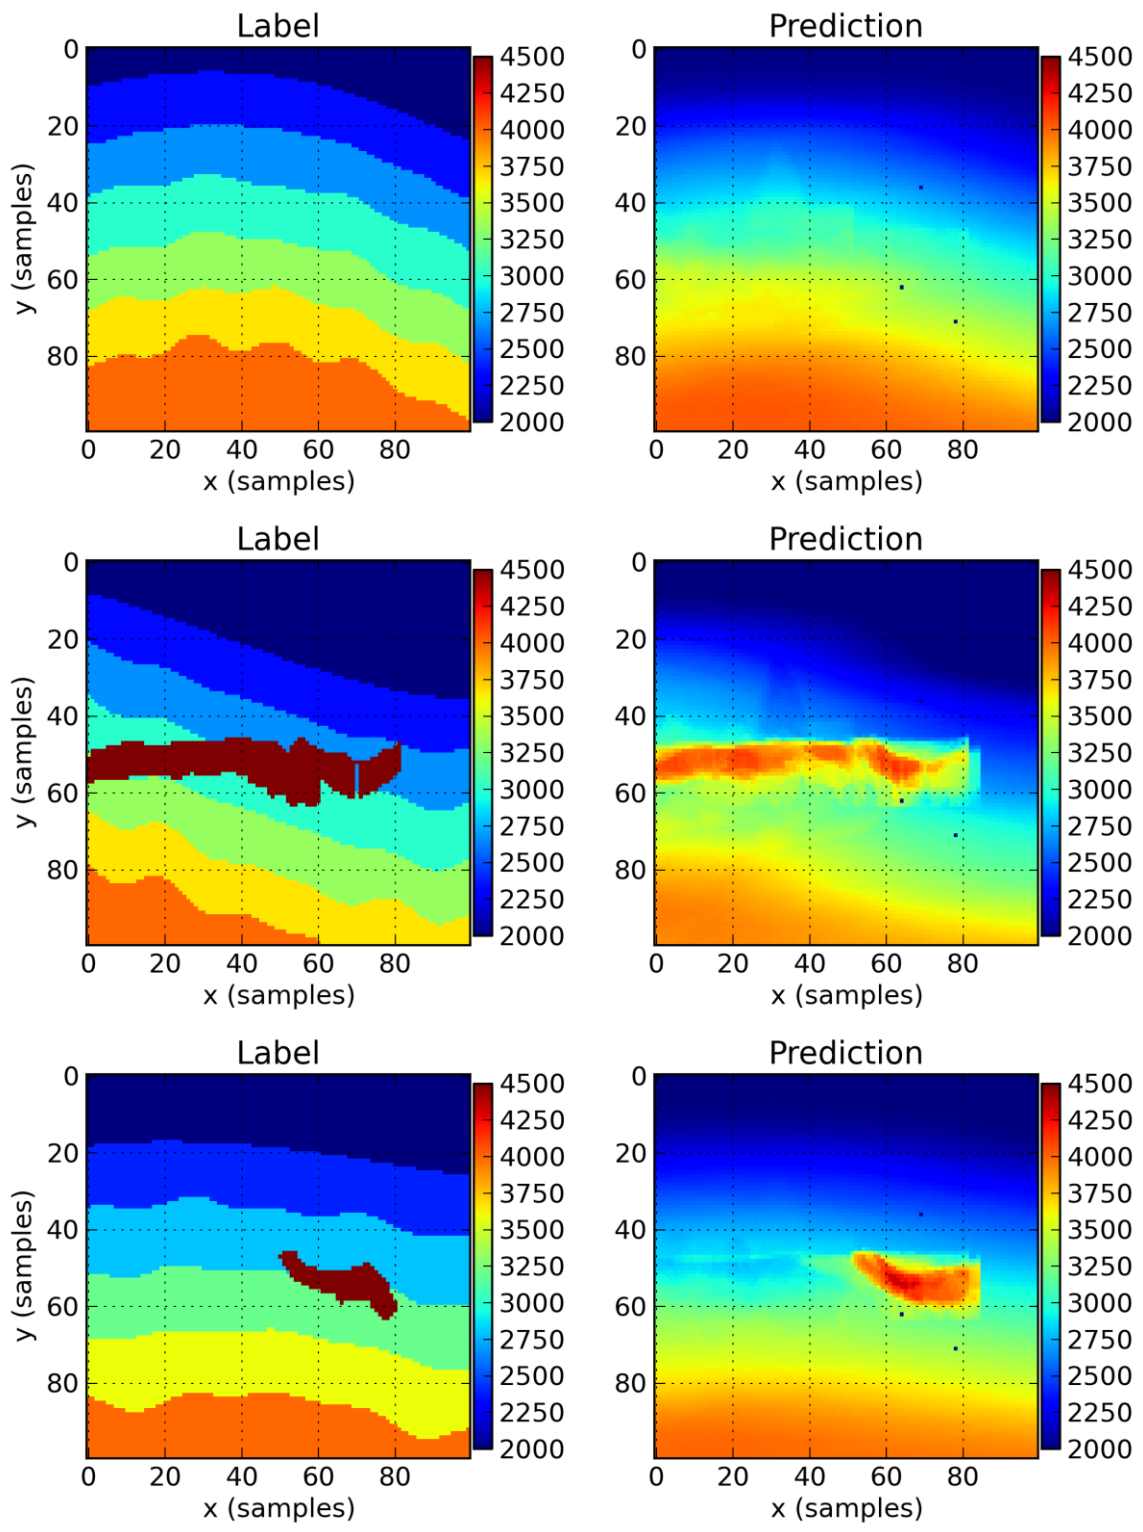


Figure 8. In the top row left column a layered model is presented as ground truth (label), the right column shows the prediction generated with the trained model. That model was training with data that only contains layered models with different number of layers and velocity per layer. The predicted model closely resembles the

label in structure and actual velocity. The mid and bottom row are results for a different trained model, this one has been trained with a dataset that also contains salt bodies, which has been handcrafted.

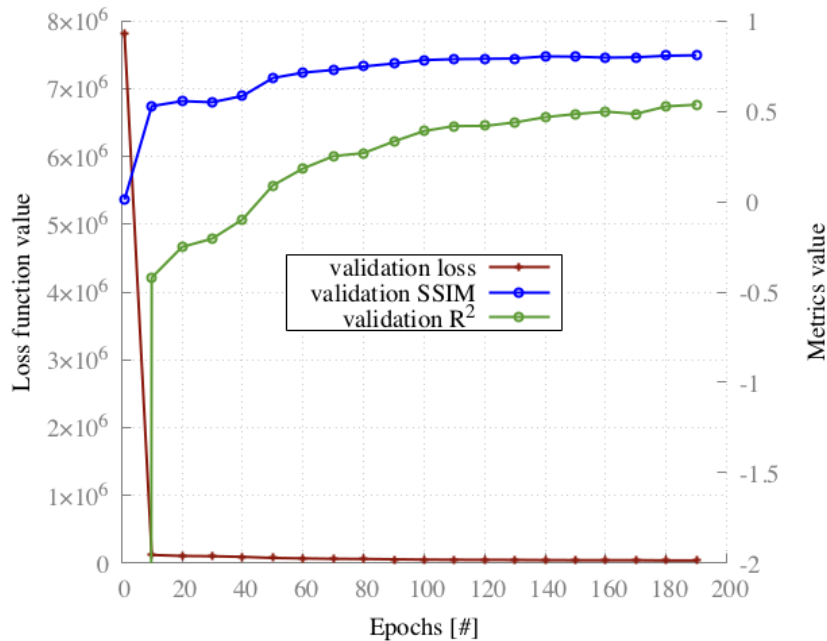


Figure 9. Validation loss function value (red curve related to left Y axis) informs about the effectiveness of the training process. Green and blue curves (related to the right Y axis) represent the evolution of the accuracy metrics during training with salt bodies. The horizontal axis represents training epochs.

In terms of prediction accuracy, for experiments that only contain layers (Figure 8 top), the R² score is 0.8124 and SSIM is 0.8939, which is comparable to the results obtained for the set of experiments of type I. For the experiments with and without salt bodies (Figure 8, mid and bottom) the R² score is 0.5536 and the SSIM is 0.8101. As expected, the task of predicting a model with salt bodies is more difficult and therefore the performance is lower than the task of predicting plain velocity models. Also, the variability of the salt bodies shape and location is more difficult to learn with the size of training dataset, that used, this explains why the R² score is more affected (Figure 9) than the SSIM for this case. In any case, the overall performance trend is positive, the salt bodies are located properly and the surrounding structure resembles the labels in direction and velocity value.

Conclusions

The concept introduced here has enormous potential, just from the computing perspective the results were computed in a single computing node, not in a cluster, further the input to the ML system has only one type of feature and the number of training samples is tens of thousands, which barely falls in the big data category. All these elements when scaled up can contribute to improve the accuracy and generalization of the predicting models.

In terms of performance, the results clearly show that most major elements of the expected models are identified and placed in their corresponding location. The results are stable and reliable as the R^2 score indicates, also they are accurate as the SSIM metric computed for all the experiments is above 0.8 (out of 1.0).

The future work spawn to three main directions: extend the method to three-dimensional data, which from the Machine Learning perspective is mainly a scalability problem rather than a fundamental one. The second line of work must deal with testing our method with real data; to that end the most sensitive part is the feature extraction step. Finally, our method can be combined with full waveform inversion, where our models can play the role of an accelerator.

Acknowledgments

The authors thank Shell Intl Exploration & Production Inc. and MIT for permission to publish this paper.

Corresponding author: Mauricio.araya@shell.com

References

Addison, V., 2016, Artificial Intelligence Takes Shape in Oil and Gas Sector: EPmag, April.

Adler, A., Boubilil, D., Zibulevsky, M., 2017, Block-based Compressed Sensing of Images via Deep Learning, The 19th IEEE International Workshop on Multimedia Signal Processing, October 16-18, 2017, London, UK.

Araya-Polo, M., Dahlke, T., Frogner, C., Zhang, C., Poggio, T., Hohl, D., 2017, [Automated fault detection without seismic processing](#), March, pp. 18-24, The Leading Edge, SEG.

Biondi, B., 2006, 3D Seismic imaging, SEG.

Burger, H. C., Schuler, C. J., Harmeling, S., 2012, Image denoising: Can plain neural networks compete with BM3D?, IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Bougher, B., and Hermann, F., 2016, AVA classification as an unsupervised machine-learning problem: In Proceedings of SEG's International Exposition and 86th Annual Meeting in Dallas, Texas.

Dahlke, T., Araya-Polo, M., Zhang, C., Frogner, C., 2016, Predicting geological features in 3D Seismic Data: In proceedings of Advances in Neural Information Processing Systems (NIPS) **29**, 3D Deep Learning Workshop.

De, S., Deo, D., Sankaranarayanan, G., Arikatla, V. S., 2011, A physics-driven neural networks-based simulation system (phynness) for multimodal interactive virtual environments involving nonlinear deformable objects, *Presence* 20, no. 4, 289-308.

Dong, C., Loy, C.C, He, K., Tang, C., 2016, Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38, no. 2, 295-307., IEEE.

Frogner, C., Zhang, C., Mobahi, H., Araya-Polo, M., Poggio, T.A., 2015, Learning with a Wasserstein loss: In proceedings of Advances in Neural Information Processing Systems (NIPS) **28**.

Goodfellow, I., Bengio, Y., Courville, A., 2016, *Deep Learning*, MIT Press.

Grzeszczuk, R., Terzopoulos, D., Hinton, G., 1998, NeuroAnimator: fast neural network emulation and control of physics-based models: In Proceedings of the 25th annual conference on Computer graphics and interactive techniques (SIGGRAPH '98).

Guillen, P., 2015, Supervised learning to detect salt body: In Proceedings of SEG's International Exposition and 85th Annual Meeting in New Orleans, Louisiana.

Hastie, T., Tibshirani, R., Friedman, J, 2001, *The Elements of Statistical Learning*, New York, NY, USA, Springer New York Inc.

Hale, D., 2013, Methods to compute fault images, extract fault surfaces, and estimate fault throws from 3d seismic images: *GEOPHYSICS*, 78, no. 2, 33-43.

Hale, D., 2012, Fault surfaces and fault throws from 3d seismic images: In Proceedings of SEG's International Exposition and 82th Annual Meeting in Denver, Colorado.

Hornik, K., Stinchcombe, M., White, H., 1989, Multilayer feedforward networks are universal approximators: *Neural Networks*, Volume 2, Issue 5, 359-366, ISSN 0893-6080, [http://dx.doi.org/10.1016/0893-6080\(89\)90020-8](http://dx.doi.org/10.1016/0893-6080(89)90020-8).

LeCun, Y., Bengio Y., and Hinton G., 2015, Deep learning, *Nature*, **521**, 436-444.

Luo, S., and Hale, D., 2012, Velocity analysis using weighted semblance: *Geophysics* 77, U15-U22

Nath, SK., Chakroborty, S., Singh, SK., Ganguly, N., 1999, Velocity inversion in cross-hole seismic tomography by counter-propagation neural network, genetic algorithm

and evolutionary programming techniques: *Geophysical Journal International*, **135**, 671–680.

Roth, G., and Tarantola, A., 1994, Neural networks and inversion of seismic data: *J. Geophys. Res.*, **99**, 6753–6768.

Szeliski, R., 2010, *Computer Vision: Algorithms and Applications*, Springer Science & Business Media.

Van der Baan, M., Jutten, C., 2000, Neural networks in geophysical applications: *Geophysics* **65**, 1032-1047.

Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004, "Image quality assessment: from error visibility to structural similarity," in *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600-612, April 2004.

Wang G., 2016, A Perspective on Deep Imaging, *IEEE Access*, vol. 4, pp. 8914-8924.

Würfl, T., Ghesu, F. C., Christlein, V., Maier, A., 2016, Deep learning computed tomography, In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 432-440, Springer International Publishing.

Xie, J., Xu, L., Chen, E., 2012, Image denoising and inpainting with deep neural networks, *Advances in Neural Information Processing Systems (NIPS)* **25**.

Yilmaz, O., 2001, *Seismic data analysis: Soc. of Expl. Geophys.*

Zhang, C., Frogner, C., Araya-Polo, M., Hohl, D., 2014, Machine-learning based automated fault detection in seismic traces: In *Proceedings of 76th EAGE Conference and Exhibition*.

BizTech, 2014, High Performance Computing' s Role in Energy Exploration, *BizTech magazine*.