AEGC2019 Data to Discovery
Australasian Exploration Geoscience Conference
2-5 September 2019 • Perth, Western Australia
Incorporating the AIG, ASEG, PESA, and WABS

Using machine learning to interpret 3D airborne electromagnetic inversions

Eldad Haber
CGI / UBC
1623 W 2nd Avenue
Vancouver, BC, CAN
haber@eoas.ubc.ca

Jen Fohring
CGI
1623 W 2nd Avenue
Vancouver, BC, CANADA
jen@compgeoinc.com

Mike McMillan*
CGI
1623 W 2nd Avenue
Vancouver, BC, CANADA
mike@compgeoinc.com

Justin Granek
CGI
1623 W 2nd Avenue
Vancouver, BC, CANADA
justin@compgeoinc.com

SUMMARY

Although 3D airborne electromagnetic inversions have improved greatly in recent years, the presence of smooth boundaries has often been a strong criticism. This smoothness can easily be remedied by applying different types of regularization and constraints to the model, but another approach is to learn what underlying structures or boundaries these smooth transitions represent.

To perform this advanced inversion interpretation, we trained a machine learning algorithm known as VNet to identify the relationship between a true synthetic model and the resulting smooth 3D inversion model. By training on one section of the model and predicting on another, the algorithm was able to learn the general relationships required to intelligently sharpen the inversion model in the prediction area. The resulting images approximate the true synthetic model to a much closer degree compared to the original inversion model. The VNet was trained in two ways, one to predict a conductivity value for each pixel, and another to predict a classification unit for each pixel presuming the conductivity for each class is known. Each method performed similarly well with some minor differences, which gives the user some options depending on the scenario and how much a priori information is known.

Overall this automatic interpretation technique worked well over a synthetic model, and future simulations will be run in order to make the method more robust and applicable for field scenarios.

Key words: machine learning, artificial intelligence, electromagnetics, inversion, interpretation.

INTRODUCTION

The inversion of airborne electromagnetic (AEM) data in 3D is continually evolving and improving with each passing year. However, the diffusive nature of airborne electromagnetic surveys makes the sensitivity to conductivity contrasts at depth much weaker compared to the surface. A typical L2 type regularization will often smooth these discontinuities at depth and to some degree near the surface, while sharper L1 or L0 regularizations may place the boundaries in the wrong place or produce spurious jumps altogether. In areas such as Australia where imaging these subtle boundaries is important for water mapping and other applications, we explore using deep-learning techniques to extract the most value out of 3D inversions.

Deep learning is a type of artificial intelligence, and multi-layer convolutional neural networks (CNNs) is a subset of deep-learning that integrates spatial data sets and extracts subtle relationships between them. CNNs have successfully been applied to mineral prospectivity mapping (Brown et al., 2000, Cracknell 2014; Granek et al., 2016; Granek and Haber, 2016) as well as water exploration using Landsat data (Isikdogan, 2017; Long, 2017).

In this work, we examine the potential of CNNs for extracting additional information from a 3D inversion model by training a VNet CNN architecture to understand the transformation between the true model and the resulting inversion model using a projected misfit calculation (Haber, 2019). Once this CNN is trained, it can transform the inversion model back into something that looks more like the true model. The thought is by training many realizations of synthetic models, field inversions can be run through this procedure to extract subtle boundaries.

Our results show that the VNet shows good promise in extracting conductivity boundaries from synthetic 3D AEM inversions, which suggests that with more work deep learning can be a valuable tool for the interpretation of inversion models.

METHOD AND RESULTS

The problem of training an algorithm for inversion interpretation can be formulated as follows: given a series of conductivity models, $\mathbf{X}^{\mathbf{train}}$, a map of true conductivities, $\mathbf{Y}^{\mathbf{train}}$, (also referred to as labels) and a function, $\mathbf{f}(\boldsymbol{\theta}; \mathbf{X}^{\mathbf{train}})$, find parameters, $\boldsymbol{\theta}$ such that

$$\boldsymbol{f}\left(\boldsymbol{\theta}, \mathbf{X}^{\mathbf{train}}\right) = \mathbf{Y}^{\mathbf{train}} \quad (1)$$

where the function, **f**, is a deep CNN known as a VNet**.** The parameters, **θ,** are made from convolution kernels, **K**, biases, **b** and linear weights, **W**.

CNNs are a class of machine learning algorithms which can model complex non-linear relationships between input and output data and are commonly applied to visual imagery and data that is spatially dependant.

A VNet architecture, as shown in Figure 1, is comprised of many convolutional levels, where the data is down-sampled to obtain lower resolution, and then up-sampled to eventually return to the original resolution. The name VNet simply refers to the V-shape of this process of down and up-sampling. It is a powerful tool that has successfully been developed and applied to medical imaging and semantic segmentation problems (Ronneberger, 2015; Milletari, 2016).

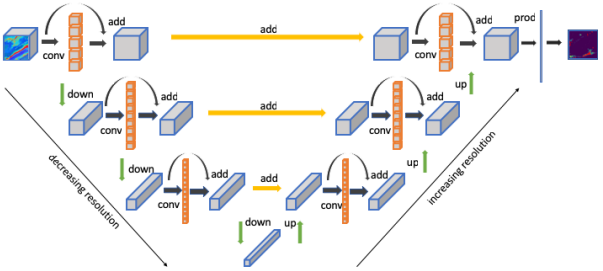

Figure 1. Graphic representation of the forward V cycle process of a VNet architecture.

A VNet can be thought of a process that constructs many images with different resolutions. The propagation of the input data, $\mathbf{X}^{\mathbf{train}}$, through a VNet combines all the features from the images at many levels of resolution. This combination is important because it allows the network to identify and learn features at many different spatial scales. This process is known to be very efficient both from a computational point of view and from avoiding local minima in the optimization (learning) process (Haber and Modersitzki, 2006).

The final level of our VNet consists of a fully connected layer, where a linear classifier is applied. This means a score is assigned to each possible data category by combining (multiplying) the propagated features with weights, **W**, to output $\mathbf{Y}^{\mathbf{pred}}$.

The "learning" is accomplished by solving an optimization problem that minimizes the loss between predicted outputs $\mathbf{Y}^{\mathbf{pred}}$, and the labels $\mathbf{Y}^{\mathbf{train}}$. This results in the estimation of the parameters of the VNet forward propagation (**K**, **W, b**), so that the CNN accurately approximates the relationship between the given data-labels.

For interpreting the 3D AEM inversion model, we formulate the deep learning problem using the following projected misfit function:

$$\min_{\boldsymbol{K},\boldsymbol{W}} ||\mathbf{P}(\,\mathbf{f}(\mathbf{K},\mathbf{W},\mathbf{b},\mathbf{X}^{\mathbf{train}}) - \mathbf{Y}^{\mathbf{train}})||_F^2 \qquad (2)$$

where $\mathbf{f}(\mathbf{K}, \mathbf{W}, \mathbf{b}, \mathbf{X})$ denotes the output of the VNet forward propagation, $\mathbf{Y}^{\mathbf{train}}$ denotes the given training labels, **P** is a projection operator that selects a subset of points in the sample grid at each iteration, and $\|.\|_F$ denotes the Frobenius norm.

The projected misfit-function in Equation 2 enables training on a subset of points, which is particularly useful for sparse labels, such as boreholes locations or for training on part of an inversion model.

To solve the optimization problem, we use the stochastic gradient descent method, which is widely used in the deep learning community. Once the parameters **K, b,** and **W** are learned, we use them to interpret parts of the inversion model that have not been trained on by simply forward propagating the input data through the VNet.

For our synthetic test example, we construct a 3D conductivity model with variably dipping layers and a near-vertical contact in the center as shown in Figure 2. The two cross-sections in Figure 2 are 4000 meters long and 400 meters in depth.

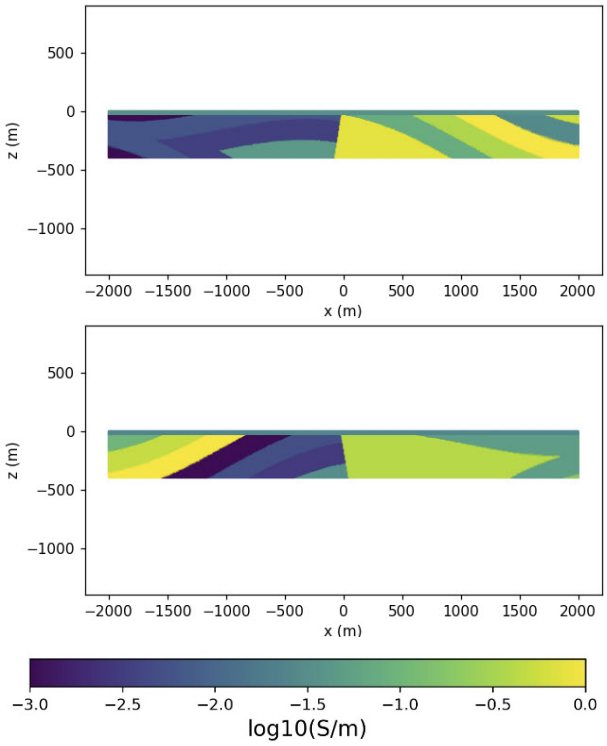

Figure 2. Top) East-West cross-section through the true model in the Southern training region (Slice 5). Bottom) East-West cross-section through the true model in the Northern prediction region (Slice 205).

The slices are separated by 400 m in the Northing direction, which demonstrates how the conductivity structures change dramatically from South to North in this synthetic example. A Skytem AEM survey (Sørensen et al., 2004) was simulated over the region with data points computed every 12 meters along-line with a 100 m line spacing. In total, five lines were simulated, and the data were contaminated with 3% Gaussian errors prior to being inverted in 3D with in-house software.

The VNet took the synthetic and 3D inversion model results and split the models into 210 east-west slices, each separated by 2m in the northing direction. The algorithm proceeded to take the southern 150 slices and learned the relationship between the blurred inversion model results and the true model. It then predicted what the true conductivities would be in the Northern 60 slices of the 3D inversion model. The slices through the true model used for the predictions are shown in Figure 3 with the corresponding 3D inversion results displayed in Figure 4.

The VNet conductivity predictions of the northern region are depicted in Figure 5 which show extremely encouraging results. The network predictions are a sharper version of the inversion model slices, and they resemble the true synthetic model to a greater degree. Slice 150 as shown in Figure 5 is the closest to the training region and thus has the best results, but even at slice 210 which is 120 m to the north of the training region, the image has a closer resemblance to the true model compared to the inversion result. Predictably, near the bottom and edges of the model is where the worst predictions are.

Figure 3. True model slices in the prediction zone.

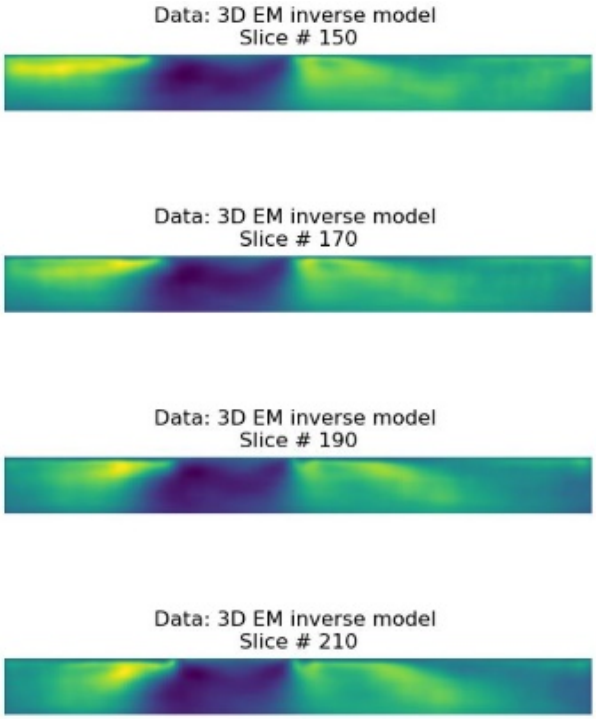

Figure 4. 3D inversion model slices in the prediction zone.

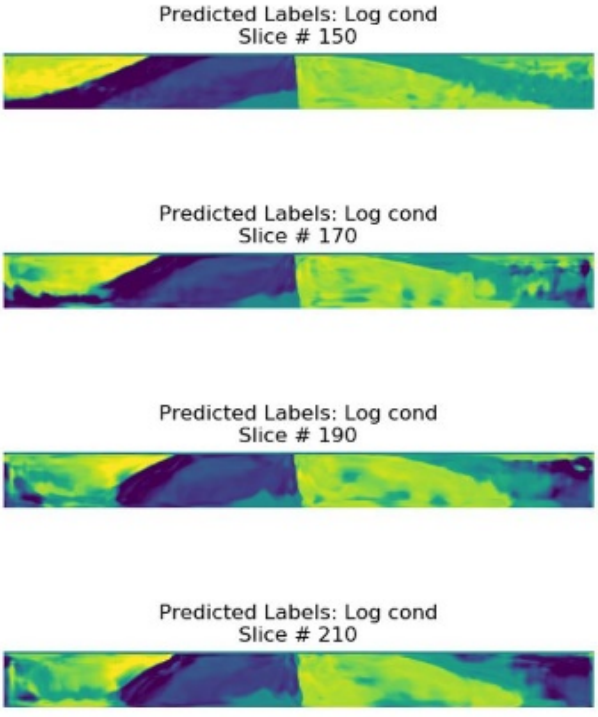

Figure 5. VNet conductivity predictions.

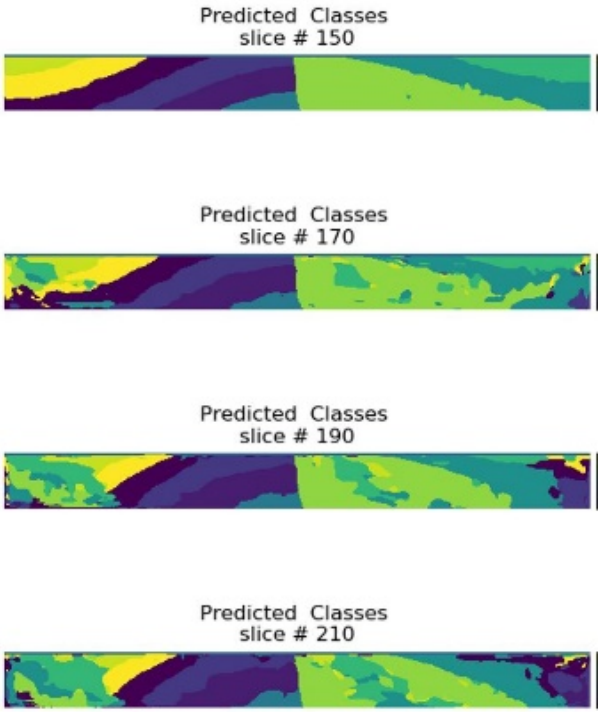

Figure 6. VNet classification predictions.

A second version of the VNet was trained to do a classification prediction. Here the network knows the true conductivity for each layer but doesn't know where the layers exist in space. For each pixel in the inversion model, the network predicts what layer it should correspond to. The classification predictions are depicted in Figure 6. Once again, the slice closest to the training section (slice 150) demonstrates the best prediction, and at slice 210 the result is able to position most of the primary boundaries in the correct location. The results from Figure 5 and 6 collectively suggest that using VNet to interpret 3D AEM inversion models is possible, but many further simulations need to be run in order to train the network for use on field data.

CONCLUSIONS

In this work we have used deep neural network and in particular, the VNet architecture to interpret 3D airborne electromagnetic inversions. The method is based on obtaining training sets, which involves a true model and an inverted one. The networks successfully learned the connection between the inverted and the true model and corrected the models to obtain more realistic and geologically interpretable models.

While the approach is promising there are a number of limitations. First, the method is not generic. It requires training models that are have similar features to the ones we need to recover. One cannot train on resistive areas and recover models on conductive areas. Obtaining geological models in general is not easy and requires more research. Nonetheless, the success of our approach encourages us to use the technique on more models and further test its utility.

REFERENCES

Brown, W.M, Gedeon,T.D., Groves, D. I, and Barnes, R.G., 2000, Artificial neural networks: A new method for mineral prospectivity mapping. Australian Journal of Earth Sciences, 47:4, 757-770.

Cracknell, M. and Reading, A.M, 2013, Geological mapping using remote sensing data: A comparison of five machine learning algorithms, their response to variations in the spatial distribution of training data and the use of explicit spatial information. Computers and Geosciences, 63, 22-33.

Granek, J. and Haber, E., 2016, Advanced geoscience targeting via focused machine learning applied to the QUEST project dataset, British Columbia. Geoscience BC Summary of Activities 2015. Geoscience BC Report 2016-1, 77-126.

Granek, J., Haber, E., and Holtham, E., 2016, Resource Management through Machine Learning. In 25th Geophysical Conference and Exhibition hosted by ASEG-PESA-AIG 2016, 1-5.

Granek, J., 2016, Application of Machine Learning Algorithms to Mineral Prospectivity Mapping, Ph.D. Thesis, The University of British Columbia.

Haber, E., and Modersitzki, J., 2006. Intensity gradient based registration and fusion of multi-modal images. International Conference on Medical Image Computing and Computer-Assisted Intervention, 726-733.

Haber, E., 2019, A deep learning geoscience package, Publication in Process.

Isikdogan, F., Bovik, A., and Passalacqua, P., 2017, Surface Water Mapping by Deep Learning. IEEE Journal of Selected Topic Applied Earth Observations and Remote Sensing, 10-11, 4909-4918.

Long,Y et al., 2017, Convolutional Neural Networks for Water Body Extraction from Landsat Imagery, International Journal of Computational Intelligence and Applications. 55, 2486-2498.

Milletari, F., Naveb, N., Ahmadi,S. 2016, V-Net: Fully Convolutional Neural Network for Volumetric Medical Image Segmentation, 2016 Fourth International Conference on 3D Vision IEEE, 565-571.

Ronneberger, O., Fischer, P., and Brox, T., 2015, U-net: Convolutional networks for biomedical image segmentation. International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 234-241.

Sørensen, K., and Auken, E., 2004, SkyTEM–a new high-resolution helicopter transient electromagnetic system. Exploration Geophysics, 35, 194-202.

Voss, K., Swenson, S., and Rodell K., 2015, Quantifying renewable groundwater stress with GRACE. Water Resources Research, 51, 1-22