Mineral prospectivity mapping using a VNet convolutional neural network

Michael McMillan[1], Eldad Haber[2], Bas Peters[1,3], and Jennifer Fohring[1]

https://doi.org/10.1190/tle40020099.1

Abstract

Major mineral discoveries have declined in recent decades, and the natural resource industry is in the process of adapting and incorporating novel technologies such as machine learning and artificial intelligence to help guide the next generation of exploration. One such development is an artificial intelligence architecture called VNet that uses deep learning and convolutional neural networks. This method is designed specifically for use with geoscience data and is suitable for a multitude of exploration applications. One such application is mineral prospectivity in which the machine is tasked with identifying the complex pattern between many layers of geoscience data and a particular commodity of interest, such as gold. The VNet algorithm is designed to recognize patterns at different spatial scales, which lends itself well to the mineral prospectivity problem of there often being local and regional trends that affect where mineralization occurs. We test this approach on an orogenic gold greenstone belt setting in the Canadian Arctic where the algorithm uses gold values from sparse drill holes for training purposes to predict gold mineralization elsewhere in the region. The prospectivity results highlight new target areas, and one such target was followed up with a direct-current induced polarization survey. A chargeability anomaly was discovered wherein the VNet had predicted gold mineralization, and subsequent drilling encountered a 6 g/t Au intercept within 10 m of drilling that averaged more than 1.0 g/t Au. Although most of the prospectivity targets generated from VNet were not drill tested, this first intercept helps validate the approach. We believe this method can help maximize the use of existing geoscience data for successful and efficient exploration programs in the future.

Introduction

The mining industry is increasingly leaning toward using advanced technologies to assist in the exploration process. The overall goal is simply to alter the continuing trend that sees a lack of new major mineral discoveries despite increased expenditures and field efforts (Barnett and Williams, 2006). As such, the mining industry is actively pursuing alternatives to traditional exploration approaches. One such alternative is to analyze the vast amount of geoscience data previously collected using artificial intelligence (AI) and, in particular, machine learning (ML) algorithms to integrate all relevant data sets. For example, in prospectivity mapping, the algorithm learns from the input geoscience data and produces a mineralization prediction map for a commodity of choice. The recent explosion of big data and the emergence of AI has created much excitement in the media; however, ML is well suited to tackle the problem of mineral prospectivity. Given that the relationship between a certain commodity and hundreds of variably noisy geoscience data sets is generally subtle and perhaps unknown, developing a computer-derived relationship without introducing a human bias is increasingly valuable.

While AI is now commonly used in fields such as computer vision with applications in facial recognition, target acquisition, satellite imaging, and threat detection (Lecun et al., 2015), there is a fundamental difference between the use of AI in these fields and its use in geoscience. Take for example the task of face detection. If we show a labeled face (i.e., a face belonging to a known person) to people who know that individual, many or all will be able to recognize the person in the image. Thus, the level of "expertise" to label the image is extremely low. In contrast, take the task of detecting potential gold mineralization from magnetic data. Many industry-leading experts can disagree on the best target locations within the same data, and, even worse, in many cases it is impossible to know which of the experts is right. The issue with labels and their uncertainty is one major difference between AI in other disciplines and AI in the natural resource industry. Accordingly, AI technology must be used slightly differently for geoscience applications.

In this study, we introduce a customized ML algorithm called VNet that learns the relationship between any type of spatial input (i.e., geoscience data) and a different spatial output (i.e., gold values from sparse drilling). We present the background to our methodology, the algorithm itself, and then show how it can be trained to obtain geologically feasible results that can be verified. To this end, we train the VNet to learn the relationship between various geoscience data sets and gold mineralization in an orogenic greenstone belt in the Canadian Arctic. Once this relationship has been learned, the network makes predictions on areas with similar geoscience data but no drilling to assess which areas have the highest potential to host undrilled gold mineralization. The results are validated on a nearby region where drill holes were not used in the training to ground truth the method. Over the primary study area, some new targets established by the VNet predictions were followed up with a direct-current induced polarization (DCIP) survey, and a subsequent drill hole over a VNet and IP target intercepted buried gold mineralization.

[1]Computational Geosciences Inc., Vancouver, British Columbia, Canada. E-mail: mike@compgeoinc.com; jenn@compgeoinc.com.
[2]University of British Columbia, Department of Earth, Ocean and Atmospheric Sciences, Vancouver, British Columbia, Canada. E-mail: ehaber@eoas.ubc.ca.
[3]Emory University, Atlanta, Georgia, USA. E-mail: bas.peters@emory.edu.

Mineral prospectivity and ML background

Initially developed in the 1980s (Bonham-Carter et al., 1989), mineral prospectivity mapping is a method that aims to derive a map of mineral potential by integrating multiple geoscience data sets, as shown in Figure 1. The typical formulation uses a set of data and an independent set of known targets. The data can be any geoscience data — e.g., geochemistry, geophysics, geology, remote sensing, etc. — and the targets should mark the place of known or thought-to-be-known mineralization. In ML language, the input geoscience layers are called training data, while the targets are referred to as labels. For most applications (but not all), the method assumes a spatial match between the training data and the labels. That is, the maps of both training data and labels are over the same spatial region, and often with similar data densities. In prospectivity mapping, the goal is to learn the relationship between data and labels to apply this newly found knowledge elsewhere that similar data exist but that labels are unavailable. In other words, where can I drill next?

Over the years, a variety of algorithms have been implemented for prospectivity mapping, such as weights of evidence (Agterberg et al., 1990), logistic regression (Harris and Pan, 1999), and, more recently, deep neural networks (Brown et al., 2000; Cracknell and Reading, 2014; Granek, 2016; Granek et al., 2016; McMillan et al., 2019). In each one, the general framework can be explained as follows. Assume there exists some collection of colocated geoscience data sets, **X**train, and for every position there also exists a small subset of points with a known value for mineral prospectivity, **Y**train. These known values can be assay results from hand samples, gold grades from drilling, or a simple binary "yes" indicating there is mineralization or a mine at that location. The mineral prospectivity algorithm is then tasked with learning the mapping function, f, such that

$$f(\mathbf{\Theta}; \mathbf{X}) = \mathbf{Y} \tag{1}$$

for all pairs **X**train and **Y**train, with trainable parameters $\mathbf{\Theta}$. The function f can take many forms from linear regression, support vector machines, decision trees, and random forests to deep neural networks. The training process is performed by demanding that the mapping can (approximately) predict the known labels and thus the known targets. To this end, we consider a loss function $\ell(\mathbf{Y}; \mathbf{Y}\text{train})$ that measures the error between the prediction error on the known data set. Common loss functions used are mean square error that sum over the (square of the) point-wise difference between the observed and predicted targets, and cross entropy when training labels have categorical values such as 0 or 1. In the training phase, we approximately minimize the loss function shown in equation 2 to obtain the optimal parameters $\mathbf{\Theta}$:

$$\min_{\Theta} \ell\big(f(\mathbf{\Theta}; \mathbf{X}\text{train}); \mathbf{Y}\text{train}\big). \tag{2}$$

With the parameter optimization completed and the mapping function learned, the user can predict mineralization elsewhere in areas with similar geoscience data but no labels (drilling).

An important issue when training any AI system is the question of validation. In many cases, solving the optimization problem in equation 2 to a high level of accuracy leads to an improvement over the training set but to a worsening overall performance. To this end, the training set is typically divided into two groups: a training set and a validation set. While we optimize the parameters over the training set, we test the performance of the function on the validation labels. The training is stopped when the validation error starts to consistently increase, indicating that additional training is starting to make the prediction worse.

While the above description of an AI system describes virtually all supervised learning methods, an important distinction between two different classes of AI algorithms should be made. The first class consists of algorithms that are geospatially unaware. This family includes almost every classical method such as weights of evidence, support vector machines, and random forests. In such algorithms, the input data are the physical, chemical, or geologic attributes of a certain point, i.e., at a specific location. Nonetheless, they do not consider the actual location of the data and surrounding spatial patterns as part of the input parameters. The second class contains algorithms that use deep convolutional neural network (CNN) architectures (Lecun et al., 2015) and represents methods at the forefront of computer vision, medical imaging, and even seismic interpretation (Ronneberger et al., 2015; Milletari et al., 2016). These algorithms aim to achieve the same goal as the first class; however, these methods treat the data as a map, or image, and use convolutions of the different attributes to learn the underlying relationship. In simplistic terms, a convolution allows the computer to not only incorporate the data at an exact location but also the surrounding patterns in the data. For example, whether a location has precisely a

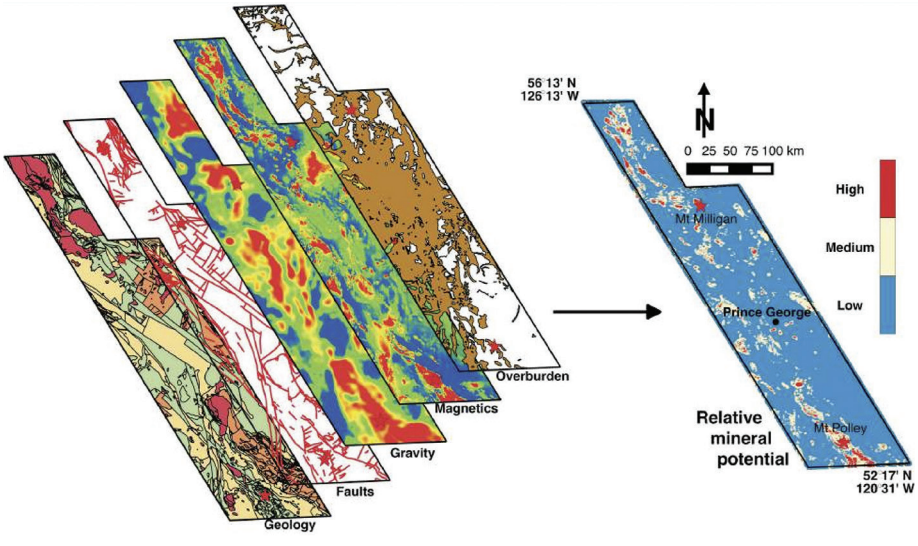

Figure 1. Schematic showing the integration of different geoscience data layers into one prospectivity map product. (From Granek and Haber, 2016).

DOI:10.1190/tle40020099.1

particular gravity value may not be as important as whether that value is in the middle of an anomalous region or residing on the edge. Given that geoscience data are geospatial and can be represented as maps or grids, there is a natural fit with a CNN approach. Furthermore, when discussing targeting with geoscience experts, they often refer to spatial patterns as a common way to identify targets. This implies that algorithms that ignore such information tend to miss known patterns and therefore can yield suboptimal results.

In this work, we introduce a deep CNN architecture known as VNet (Peters et al., 2019b), which is shown schematically in Figure 2. The VNet architecture has three main components, the first being a convolution; that is, the input data are convolved with kernels that are optimized to give different spatial representations of the data. The second component is down-sampling, which combines data at different resolutions in order to look for patterns in the input data at multiple spatial scales. The final component is interpolation from a coarse to a fine mesh. The three components together form an architecture designed to yield final predictions that combine representations of the high-resolution data at fine and coarse scales (Peters et al., 2019a). Since the data exhibit multiscale characteristics, the VNet architecture integrates and combines the data at different scales to obtain the best possible prediction.

When considering different CNN algorithms, it has been observed that deep neural networks outperform their shallower counterparts (Lecun et al., 2015), and they allow for more complex relationships. While there is no sharp distinction between the depth of networks, shallow networks generally contain very few layers, typically on the order of one to three. Recent work on deep networks, including the work here, typically uses hundreds of layers. This is important with mineral prospectivity because of the variety of data used and the complex relationship that is sought. The VNet framework is also set up to be as general as possible and can be applied to other geoscience problems such as airborne induced polarization detection, enhanced 3D inversion interpretation, and seismic horizon detection (Peters et al., 2019a).

To set up the VNet architecture for prospectivity mapping, the data, **X**train, are preprocessed to a series of images on a common grid, such as that displayed in Figure 3. The training labels, **Y**train, are gold values in drill assays for this study, which exist at sparse locations throughout the common grid. Finally, the mapping function *f* represents the entire VNet architecture, which

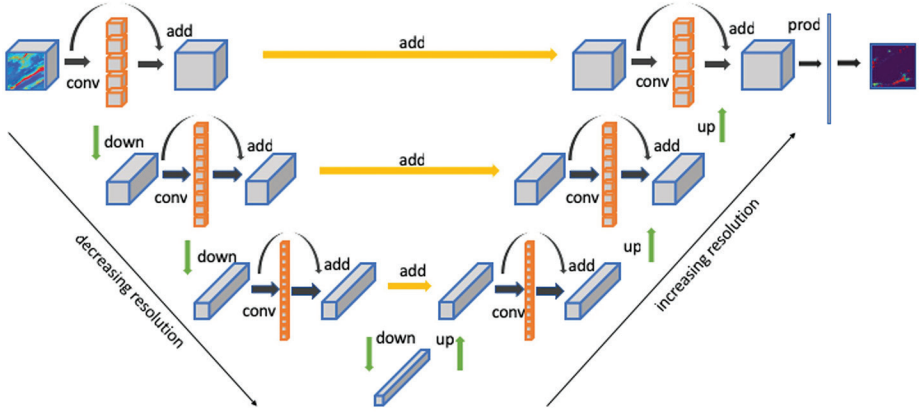

Figure 2. Graphical representation of the forward V cycle of the VNet architecture, where conv refers to convolutions.

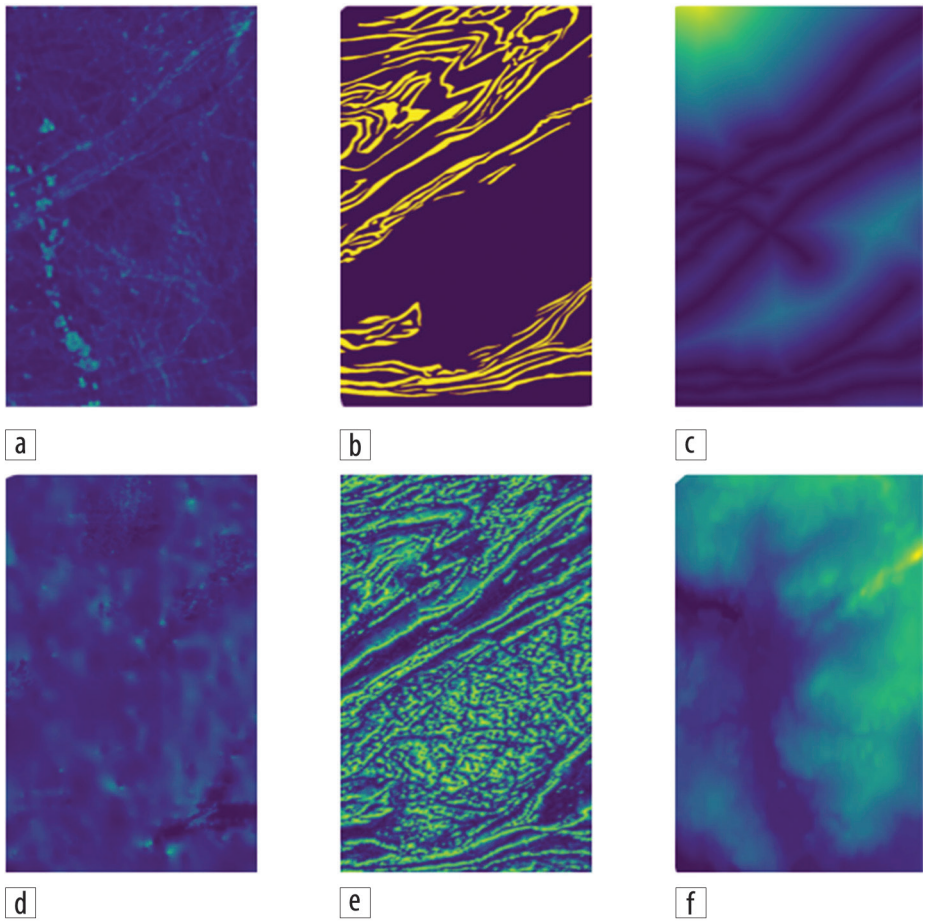

Figure 3. A selection of geoscience data layers at Committee Bay used for this study. (a) Airborne electromagnetics from a particular time channel. (b) Geologic interpretation where one lithologic unit is highlighted in yellow. (c) Distance from major faults. (d) Geochemistry from a particular element in till sampling. (e) Airborne magnetics, reduced-to-pole. (f) Topography.

includes a classifier at the end to generate a numerical gold value at each pixel in the grid. The gold values at each training and validation label location are incorporated into the loss function to evaluate how well the predicted data match the true label data. The learning process then iterates until a desired target level of fit has been achieved or the validation data misfit begins to increase over multiple iterations. Alternatively, the training labels can be separated into "mineralization" and "no mineralization" classes based on a user-defined gold grade cut-off, which adds flexibility and user options to the exercise. Regardless of the type of label used, when drilling is only available in a few locations across the survey area, a projected data misfit method is used where the data misfit is only calculated at the specific points with labels (Peters et al., 2019a). This prevents the user from having to assign and estimate label values in regions without them.

When assigning labels, an important note to remember is that the data may not be balanced. While we know where some of the targets of interest are, we do not know exclusively where they are not. Therefore, a naive choice of labels may lead to incorrect or biased training. For example, if we choose only the high gold values for the training, such that 90% of our labels show economic mineralization, a simple way for the algorithm to fit the data is to choose parameters such that the mapping function f predicts mineralization everywhere. To this end, we make sure that the labels are roughly balanced, with approximately an equal number of "good" outcomes and "bad" outcomes. This way, the training is designed to not favor one particular outcome.

Field test — Committee Bay orogenic gold exploration

For this work, several geoscience data sets were provided from the Committee Bay Project in Nunavut, Canada, by Auryn Resources (now Fury Gold Mines). The data sets include airborne magnetic data, frequency domain airborne electromagnetic data, geologic mapping, structural interpretation, topography, geochemical assays from drill holes, and geochemical till samples. Gaps in data sets are a huge topic of research because most projects do not have a uniform coverage across the area of interest. Fortunately, except for sparse drilling labels, there was roughly uniform data coverage across the region, and any gaps in data layers were filled by kriging or directional gridding.

At Committee Bay, the commodity of interest is orogenic gold mineralization, and gold assays from the sparse drill holes were used as labels. The study area consisted of two main regions: one larger area known as Three Bluffs and one smaller nearby area known as Anuri. The goal was to generate targets for an upcoming drill program in the Three Bluffs region and to use the Anuri area for validation of the technique. As most of the geologic units at Committee Bay have a near-vertical dip, a 2D prospectivity approach was used; however, this method can also be deployed in a 3D style using inversion models and other 3D data if desired. Since gold assays have spatial coordinates in 3D, the gold values were projected vertically upward to the surface to create the 2D labels. As a prospectivity output is broken up into grid cells that define the resolution of the map (20 × 20 m for this study), after the drill holes have been projected to the surface there can often be multiple gold labels that reside within a single cell. As such, there are multiple ways to generate the final gold label values. One can take the maximum gold value in each cell, the mean value, the median value, or a weighted summation of all values within the cell. For this project, the goal was to target high-grade gold mineralization, and therefore the maximum gold value was chosen. More work needs to be done to analyze the difference between the predictions with each of the various label choices. Additionally, for targeting high grades over large intervals, some sort of integrated gold metric might be a suitable label choice.

Results

We generated the gold prospectivity map at Three Bluffs by training our VNet algorithm on the available geoscience data sets and using the projected gold labels from drilling. We first validated the method by predicting gold values at the nearby Anuri area and compared our gold predictions to known gold mineralization from drilling there. Figure 4 shows the gold prediction map at Anuri overlaid on the interpreted geology with the main geologic units numbered and labeled. Only gold prediction values above 1.0 g/t Au are displayed with a heat-map color scale that linearly stretches from 1.0 (purple) to 2.0 g/t (white) Au. The VNet result primarily shows prospective areas in the northeast corner and in the center of the survey area where the bulk of the gold mineralization has been found thus far. A zoomed-in version of the core mineralized area is shown to the right in Figure 4, which further shows the generally good correlation with existing gold values and the VNet prediction.

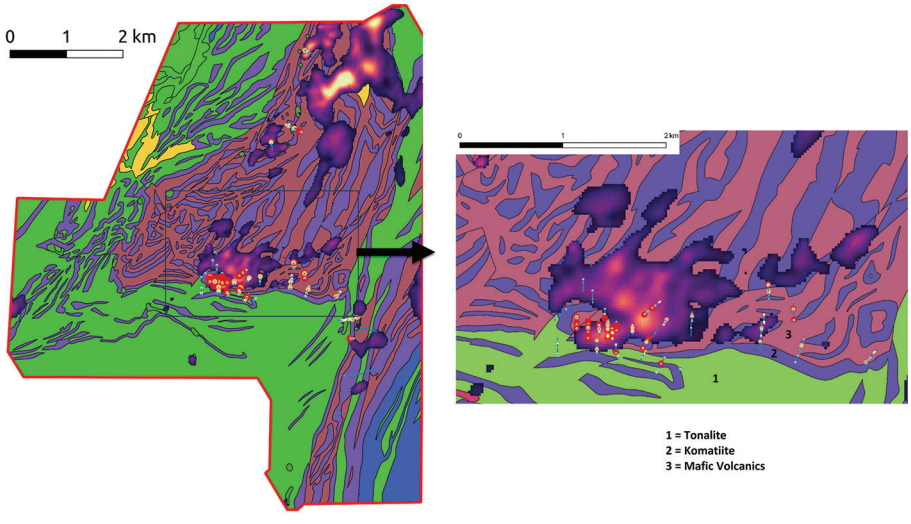

Figure 4. Validation gold prospectivity map over neighboring Anuri property, showing only prospective areas greater than 1.0 g/t Au in the clipped heat map with a zoomed-in version of the core mineralized zone. The interpreted geology map is underlain with major lithologic units numbered. Gold values from drilling are plotted as colored circles, where red represents values greater than 1.0 g/t Au.

Together this validation exercise demonstrated how the training at Three Bluffs was able to predict both known mineralization locations at Anuri and new potential areas of interest.

Since the validation test demonstrated an ability to predict known gold occurrences at Anuri, this added confidence to our results and method. Subsequently, a gold prospectivity map was generated for the Three Bluffs area. The full gold prediction map for this much larger area overlaid on the interpreted geology map is shown in Figure 5 with main geologic units numbered and labeled. Once again, only prediction zones greater than 1.0 g/t Au are displayed with a heat-map color scale that linearly stretches from 1.0 (purple) to 2.0 (white) g/t Au. Drilling values are also plotted with drill intercepts above 1.0 g/t Au in red. The two white rectangles refer to the main deposit region where the bulk of drilling is located (rectangle a) and a new target region identified by the prospectivity mapping near some existing drilling (rectangle b).

Many targets are highlighted in this exercise at Three Bluffs; however, prospective regions, or areas above 1.0 g/t Au, only represent about 1% of the total area, which in turn greatly reduces the ground to focus on. A zoomed-in version over the main deposit area and a prospective region in the north is shown in Figure 6 with a black and white color stretch for the prospectivity map, where gray and white regions are greater than 1 g/t Au. Figure 6a shows the bulk of the labels in the main deposit region along with the final training prediction. It must be noted that the prediction does not exactly match the drilling outline. In this manner, the prospectivity map acts like an inversion, where the harder the data are fit, the more artifacts emerge in the rest of the area. Therefore, there is a balance of fitting yet not overfitting the data. We help alleviate this issue by using the independent validation data at Anuri to determine when to stop the training at Three Bluffs.

Figure 6b shows an AI target in the north-central part of the survey area with no drilling, although it is next to some historic drill holes. The proximity to existing drilling made this target logistically easier than many other targets, and as such this target was chosen as the first to test. A 2D pole-dipole DCIP survey was subsequently planned, and the first line collected over the target accompanied by its 2D chargeability inversion cross section is shown in Figure 6b. The draped 2D section shows the main chargeability anomaly matches well with the outline of the AI-derived prospective zone. Figure 7 goes on to show both the resistivity and chargeability 2D inversions from this DCIP line and the drill hole that followed. The drill results show an interception with altered komatiite and disseminated sulfides through the conductive and chargeability anomaly. This anomaly coincides with the location of the prospectivity map prediction, and

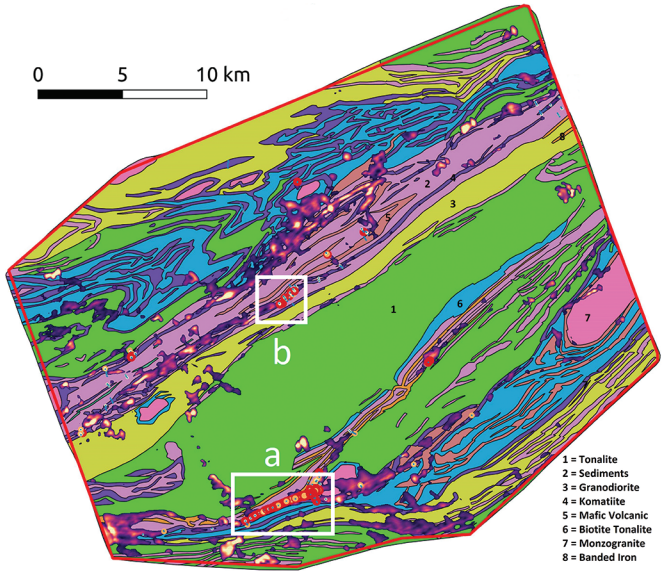

Figure 5. Gold prospectivity map at Three Bluffs, showing only the prospective areas as hot zones on the clipped heat map. The geology map is underlain with major lithologic units numbered. Gold values from drilling are plotted as colored circles, where red represents values over 1.0 g/t Au. The two white rectangles are areas of particular interest with rectangle "a" being the main deposit area and rectangle "b" representing a new target region highlighted by the AI-generated prospectivity map.

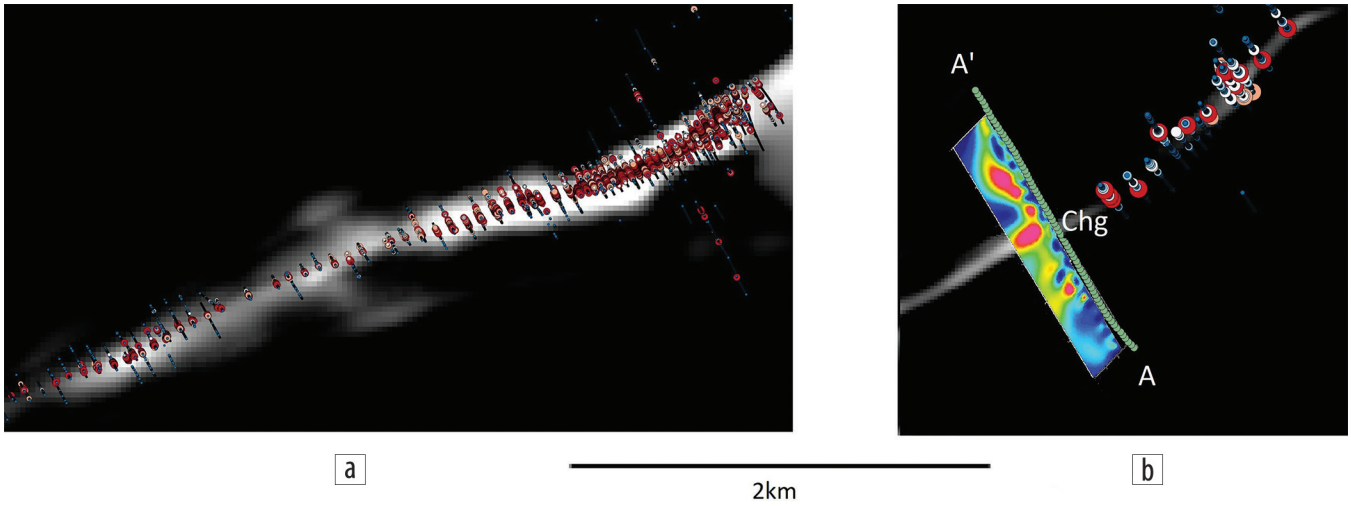

Figure 6. Zoomed-in prospectivity map in black and white for two areas at Three Bluffs. The gray and white color regions refer to areas where the prospectivity map predicted gold greater than 1.0 g/t Au. Drill holes are plotted for reference with gold grades greater than 1.0 g/t Au in red. (a) Three Bluffs main deposit region. (b) Prospective target in the north-central part of the survey near some previous drilling. Draped image is a 2D chargeability inversion cross section with station locations plotted as green dots. Red = chargeable.

DOI:10.1190/tle40020099.1

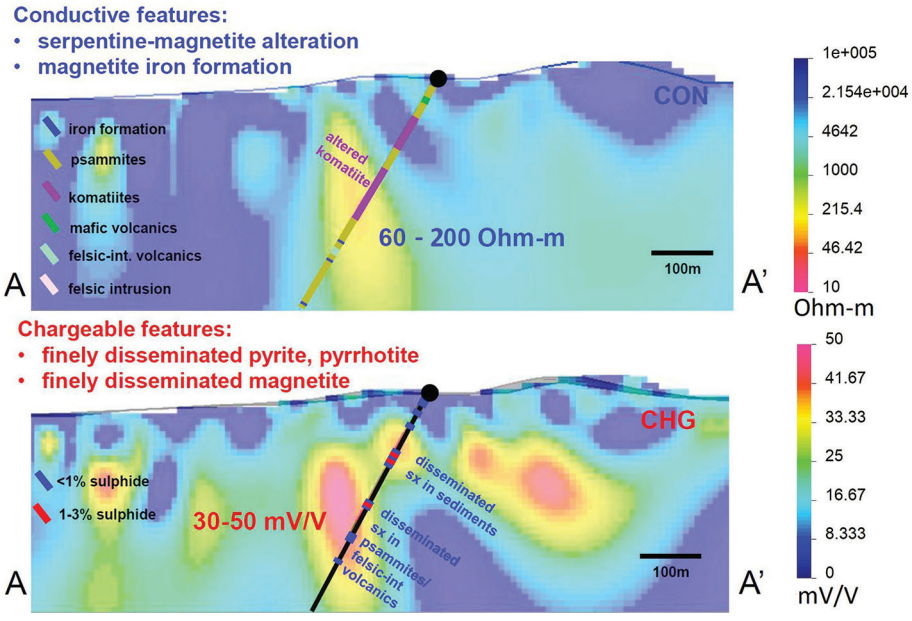

Figure 7. Drilling results overlaid on DCIP 2D inversion results from line shown in Figure 6. At top is 2D resistivity inversion, and at bottom is 2D chargeability inversion.

subsequent assay results returned interceptions of 10.5 m of 1.22 g/t Au including 1.5 m of 6.2 g/t Au in this drill hole. Although these grades are not high enough to be considered a discovery drill hole, the results demonstrate an ability for our ML platform to identify targets that host significant buried gold mineralization.

Despite the success on this hole, the inherent drawback with such an exercise is the inability to field test all the anomalous areas. Many of the targets have not been tested, and some will most likely never be tested. However, any new drill information or geoscience data collected over the area can be input back into the prospectivity algorithm to create an updated prediction. This exercise therefore should be viewed as an iterative process that improves with each year as more and more information is collected and added to the training.

Discussion and conclusions

In this work, we discussed methods for prospectivity mapping that range from primitive point-wise methods to modern algorithms based on deep CNNs. We use the latter to introduce a new method for mineral prospectivity mapping using a multi-resolution VNet architecture. The algorithm was subsequently used to integrate a large regional gold exploration data set in the Canadian Arctic to predict gold mineralization. The project included numerous geophysical, geochemical, and geologic layers, and the network was trained using gold assay values from drilling to produce a prospectivity map highlighting new exploration targets. The algorithm was first validated on a nearby region to ground truth the predictions. Here, gold predictions were made and compared to gold values from drilling not used in the training process. The results were promising, highlighting both known gold occurrences as well as new potential targets. A prospectivity map for the Three Bluffs area was then created, followed by an induced polarization survey designed to test a pronounced AI anomaly. The DCIP survey showed a conductive and chargeable anomaly that coincided with the prospectivity target, and subsequent drill testing intercepted 6 g/t Au mineralization.

Future research is ongoing to address a couple of main challenges that remain in prospectivity mapping, namely missing data and labels. Many data layers often have regions where no data have been collected, and one may still want to use the available information in the prediction. One option is to train the VNet to predict the missing data, like a prospectivity map. This AI-derived data layer can then be used as an input for the full prediction with all the data layers. A drawback of such an approach is it uses the complete data layers in both the prediction for the missing data and the ensuing prospectivity prediction. This essentially weights these inputs with more importance; however, this may be preferential to ignoring the sparse data altogether. As for preparation of labels, in this study we used the maximum gold grade in a cell as the label. One could also bin the gold grades into certain categories (such as high, medium, and low according to the grades found at the project) and then use a segmentation approach to place each location in the spatial domain into one of those categories.

Overall, there is much research left to be done to optimize the use of ML for geoscience applications, but we believe the VNet approach offers a modern option for prospectivity mapping. Because geoscientists are routinely looking for patterns in data to identify future drill targets, it is only logical to utilize the power of the machine to help accomplish the same task. Ultimately the geoscientist can vet the AI targets with their own knowledge and opinions to generate targets with an added level of confidence. TLE

Acknowledgments

The authors thank Fury Gold Mines for permission to show the prospectivity results.

Data and materials availability

Data associated with this research are confidential and cannot be released.

Corresponding author: mike@compgeoinc.com

References

Agterberg, F. P., G. F. Bonham-Carter, and D. F. Wright, 1990, Statistical pattern integration for mineral exploration, *in* G. Gaal and D. F. Merriam, eds., Computer applications in resource estimation: Prediction and assessment for metals and petroleum: Pergamon Press, 1–21, https://doi.org/10.1016/B978-0-08-037245-7.50006-8.

Barnett, C. T., and P. M. Williams, 2006, Mineral exploration using modern data mining techniques: First Break, **24**, no. 7, https://doi.org/0.3997/1365-2397.24.1097.27027.

Bonham-Carter, G. F., F. P. Agterberg, and D. F. Wright, 1989, Weights of evidence modeling: A new approach to mapping mineral potential, *in* F. P. Agterberg and G.F. Bonham-Carter, eds., Statistical applications in the earth sciences: Geological Survey of Canada, 171–183.

Brown, W. M., T. D. Gedeon, D. I. Groves, and R. G. Barnes, 2000, Artificial neural networks: A new method for mineral prospectivity mapping: Australian Journal of Earth Sciences, **47**, no. 4, 757–770, https://doi.org/10.1046/j.1440-0952.2000.00807.x.

Cracknell, M. J., and A. M. Reading, 2014, 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 & Geosciences, **63**, 22–33, https://doi.org/10.1016/j.cageo.2013.10.008.

Granek, J., 2016, Application of machine learning algorithms to mineral prospectivity mapping: PhD thesis, University of British Columbia.

Granek, J., and E. Haber, 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, 117–126.

Granek, J., E. Haber, and E. Holtham, 2016, Resource management through machine learning: 25th Geophysical Conference, ASEG, Extended Abstracts, https://doi.org/10.1071/ASEG2016ab253.

Harris, D., and G. Pan, 1999, Mineral favorability mapping: A comparison of artificial neural networks, logistic regression, and discriminant analysis: Natural Resources Research, **8**, 93–109, https://doi.org/10.1023/A:1021886501912.

LeCun, Y., Y. Bengio, and G. Hinton, 2015, Deep learning: Nature, **521**, 436–444, https://doi.org/10.1038/nature14539.

McMillan, M., J. Fohring, E. Haber, and J. Granek, 2019, Orogenic gold prospectivity mapping using machine learning: 2nd Australian Exploration Geoscience Conference, ASEG, Expanded Abstracts, https://doi.org/10.1080/22020586.2019.12073020.

Milletari, F., N. Navab, and S.-A. Ahmadi, 2016, V-net: Fully convolutional neural networks for volumetric medical image segmentation: Fourth International Conference on 3D Vision, 565–571, https://doi.org/10.1109/3DV.2016.79.

Peters, B., J. Granek, and E. Haber, 2019a, Multiresolution neural networks for tracking seismic horizons from few training images: Interpretation, **7**, no. 3, SE201–SE213, https://doi.org/10.1190/INT-2018-0225.1.

Peters, B., E. Haber, and J. Granek, 2019b, Neural networks for geophysicists and their application to seismic data interpretation: The Leading Edge, **38**, no. 7, 534–540, https://doi.org/10.1190/tle38070534.1.

Ronneberger, O., P. Fischer, and T. Brox, 2015, U-net: Convolutional networks for biomedical image segmentation: International Conference on Medical Image Computing and Computer-Assisted Intervention, 234–241, https://doi.org/10.1007/978-3-319-24574-4_28.