AEGC Australasian Exploration Geoscience Conference Brisbane 2021

Almost automatic geological mapping from AEM surveys

David Annetts
CSIRO
26 Dick Perry Av.
KENSINGTON 6151
david.annetts@csiro.au

Juerg Hauser
CSIRO
2-40 Clunies Ross St
ACTON 2601
juerg.hauser@csiro.au

SUMMARY

Interpretation of electromagnetic survey data for underlying geology is a common task that is complicated by the volume of data collected in airborne electromagnetic (AEM) surveys. We test a supervised learning approach for AEM data collected in the La Grange groundwater area, Western Australia. We use machine learning to identify the most likely geological setting at each station and use this to derive the probable extent of the seawater interface. We employ standard techniques such as cross validation to benchmark machine learning algorithms such as nearest-neighbour, naive Bayes and support vector networks.

The good agreement between a qualitative interpretation and the best-performing machine learning algorithm, here a random forest algorithm, for the seawater interface extent suggests that automatic classification has the potential to speed up the interpretation of large airborne electromagnetic surveys. Our results also suggest that careful use of machine learning algorithms trained on high quality interpretations can lead to more objective geological interpretations particularly when airborne electromagnetic data are collected in order to map regional geology. A modest effort spent interpreting small but representative survey portions can be leveraged to geological mapping of the survey as a whole.

Key words: machine learning, airborne, electromagnetic, mapping

INTRODUCTION

One reason to collect airborne electromagnetic (AEM) data is to map the near surface geology. A geological map of a survey area can be derived by determining for each station the geological structure it represents. Conventional interpretation of AEM surveys is often based on a line-by-line interpretation of the inversion results which are typically done on a station-by-station basis. Data must be inverted for a model of the earth because of the nonlinear relationship between measurements of the secondary magnetic field and electrical resistivity (or its inverse, electrical conductivity). For a large survey this can be a time consuming and often subjective process.

If we can train a machine-learning algorithm to classify data using a few carefully interpreted lines then that algorithm could be used to classify data on the remaining survey lines. Therefore, the task of geological interpretation can be sped up or at least aided by machine learning. In this study we will use the La Grange AEM survey inversion results (Annetts et al., 2016) to contrast a geological interpretation using some common machine learning algorithms with a conventional interpretation where inversion results have been analysed on a line-by-line basis. Also with specific reference to La Grange AEM survey inversion results, we discuss how a machine-learning algorithm approach might be used when inversion results are not available. We begin this paper with an overview of salient machine-learning algorithm techniques employed in classification of data.

MACHINE-LEARNING ALGORITHMS

Finding the best classifier for any problem invariably involves using cross validation to benchmark the algorithms and chose the most appropriate one for the problem at hand. It is important to keep in mind that high quality data are often more important than a superior algorithm (Caruana and Niculescu-Mizil, 2006). In other words, careful interpretation of the test and training data, and the choice of an appropriate feature vector can often be more important than identifying the best algorithm among well performing ones for the problem at hand. We use machine learning algorithms (with default parameters) as implemented in *Mathematica 11.01* (Wolfram Research Inc., 2016) in this work, and note that all algorithms used in this paper are have been implemented in a variety of programming languages. We continue with a brief review of some of the machine-learning algorithms employed in this paper.

The *k*-nearest neighbours (NN) (Cover and Hart, 1967) is a simple machine learning algorithm that employs distance (in feature space) to determine classification. Clearly, the success of such a scheme will depend on both the number of neighbours (*k*) used and the distance metric. Heuristic techniques (e.g. Hallet al., 2008; Samworth, 2012) can be used to determine an optimal number of neighbours and their weighting for a given training data set. The NN machine-learning algorithm is often used as a baseline in comparisons of machine-learning algorithms.

The Naive Bayes (NB) machine learning algorithm refers to a class of algorithms which assume that features are independent. The ‘Bayes’ part of the name refers to the probability that features are indicative of the class, while the ‘Naive’ part of the name stems from the fact that assumptions of parameter independence are often incorrect. Zhang (2004) has shown that despite the simplicity of NB schemes, they can perform quite well in practice, especially for supervised-learning problems as is the case here. Nevertheless, NB schemes are often outperformed by other algorithms (Caruana and Niculescu-Mizil, 2006; Ng and Jordan, 2001; Niculescu-Mizil and Caruana, 2005).

Like NB algorithms, the logistic regression (LR) is a linear classifier. Developed by Cox (1958), the LR algorithm can be regarded as a special case of linear regression in which probabilities between categorical dependent variables and independent variables are measured using a logistic function.

Artificial neural networks (ANN) (Bishop, 1995) attempt to solve problems in a manner similar to the way human brains are thought to solve problems. The brain's learning mechanism is typically modelled using a combination of thresholding functions in layers which can accept multiple inputs. Often, logistic functions are employed as thresholding functions. Thresholding determines whether information is passed to another layer of the ANN, and the direction of information flow can be forwards or backwards. Backwards information flow through the network is known as 'back propagation' and was a critical step in ANN development because it permitted refinement of the network based on desired outcomes. Historically, ANN's have had mixed success, though recent developments, in the form of recurrent neural networks, have been more successful in fields such as speech processing (Zen and Sak, 2015).

The Random Forest (RF) machine-learning algorithm was developed by Ho (1995) from decision trees. Random Forests (RF) proceed by constructing a multiplicity of decision trees during training. Bagging and feature randomness are used when constructing each individual tree with the aim of creating an uncorrelated forest. Given that forest predictions can be made by committee, they are more accurate than the prediction from a single tree. Breiman (2001) improved the approach of Ho (1995) by introducing bootstrap aggregating. Lin and Jeon (2006) have pointed out that a random forest classifier can be viewed as a weighted neighbour classifier with weights equal to the average of the individual trees. Some RF algorithms (e.g. Prinzie and Van den Poel, 2008) use NB and LR as integral components.

Support Vector Networks (SVN) attempt to map input vectors into a high-dimensional space in such a way that data in that space are separated using optimal hyperplanes. Vapnik (1982) showed that only a small amount of the training data (termed 'support vectors') were needed in order to determine optimal hyperplanes separating categories thus offering the possibility of an efficient algorithm after minimal training. Current implementations have their genesis in Cortes and Vapnik (1995). Zuo and Carranza (2011) describe the use of an SVN to construct maps of mineral prospectivity.

Caruana and Niculescu-Mizil (2006) evaluated the performance of all algorithms used in this paper for applicability in a range of scenarios. They found that while some methods (e.g. RF and SVN) performed better than others (e.g. NB and LR) on average, there was considerable variability across problems and metrics. Good algorithms were shown to perform poorly on particular data sets, and algorithms with generally poor performance could perform particularly well. This stresses the importance of cross validation, that is, splitting the interpreted (classified) data into training and validation sets, and measuring the performance of various machine learning algorithms in order to choose the best-performing one for the problem at hand.

THE TARGET

The La Grange AEM survey (Figure 1) was collected using a TEMPEST AEM system (Lane et al., 2000) and was designed to address a paucity of hydrogeological knowledge (Paul et al., 2013) in the La Grange groundwater area, a region of emerging importance for agriculture, minerals and petroleum exploration. The La Grange groundwater area (LGGA) lies in the western half of the onshore Canning Basin which is located in the north west of Western Australia. Although the Canning Basin has long been recognised as an important aquifer system (Forman and Wales, 1981; Lau et al., 1987; Ghassemi et al., 1992), specific knowledge of the La Grange groundwater area was sparse. Paul et al. (2013) reviewed the properties of the Broome sandstone aquifer in the La Grange groundwater area with the goal of summarising hydrogeological data, thereby establishing a baseline for future studies.

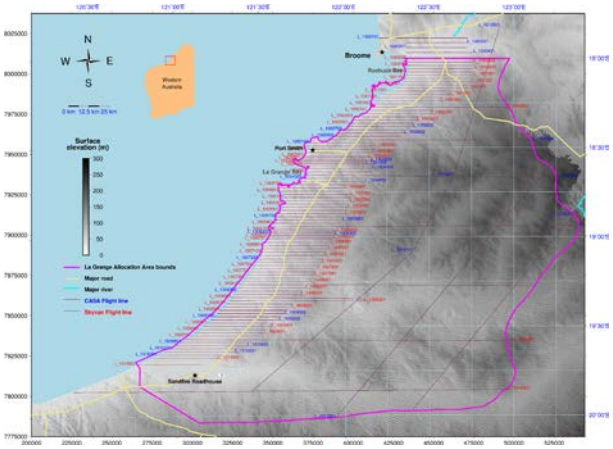

Figure 1: Location of field data used in this paper. The La Grange Groundwater Allocation Area is highlighted in magenta, and AEM survey lines are indicated with labels at their start.

The main aquifer in the La Grange groundwater area is the Broome sandstone. The Broome sandstone was formed in the lower Cretaceous period and is overlain by quaternary sediments. In the majority of the La Grange groundwater area the Broome sandstone overlies the Jarlemai siltstone which was formed in the upper Jurassic period and is an aquitard. In the south of the La Grange, the major aquifer transitions from the Broome sandstone to the Wallal sandstone. In the west of the La Grange groundwater area, the Broome sandstone is known to be affected by seawater intrusion. The seawater intrusion can extend up to 40 km inland at depth (Annetts et al., 2016; MIRA Geoscience, 2015). The Broome sandstone is an unconfined aquifer and thus knowledge of the depth to the top of of the Jarlemai siltstone (the base of the aquitard) and the depth and inland extent of the seawater intrusion over the La Grange groundwater area allows calculation of the volume available for water storage. Prime objectives of the AEM survey over the La Grange groundwater area were therefore to image the depth to the top of the Jarlemai siltstone and the spatial extent and depth to top of the seawater intrusion. The representative geology for the area is given in Figure 2. Annetts et al. (2016) inverted the AEM data for a five layered

earth using GALEISBSTEM (Brodie, 2016) and interpreted those inversion results to derive the aforementioned hydrogeological features.

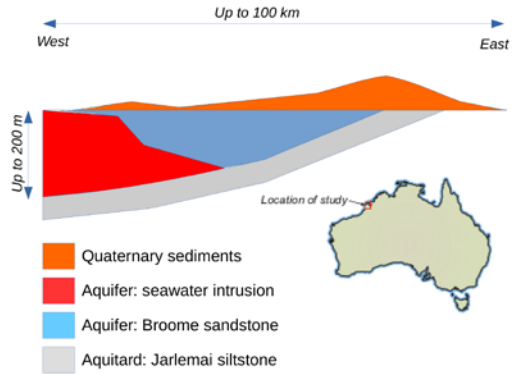

Figure 2: Cartoon illustrating geology in the LGAA. The target of interest in this paper is the extent of the seawater intrusion into the Broome Sandstone aquifer.

Critical for the success of a machine learning approach for interpretation of an AEM survey in terms of geology, is an appropriate training data set. In practice this is a geological interpretation of a subset of the stations. For supervised learning it is important that the training data set contains all the features we expect to encounter in the rest of the survey. Based on the regional geology we defined three geological sequences viz. seawater intrusion (S), quaternary sediments overlying Broome sandstone overlying seawater intrusion (QBS) and quaternary sediments overlying Broome sandstone overlying Jarlemai siltstone (QBJ). The training dataset is formed by the survey line L1006601 and Figure 3 shows the result of a station-by-station inversion for a five-layer model of the earth with the line segments that were identified as being representative of one of the three geological sequences.

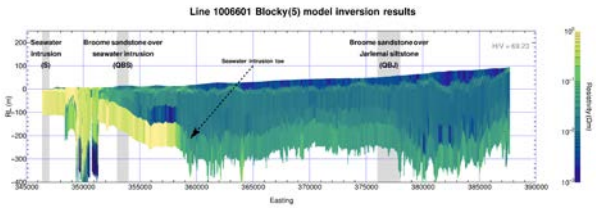

Figure 3: Inversion results from the line used to train the ML algorithms. Named features are indicated and their extent highlighted as gray backgrounds.

Figure 3 highlights the main reason to employ a machine-learning algorithm to analyse the inversion results. Because 1D models with few parameters were used, it is a relatively straightforward task to interpret geological horizons, such as the tops of the Broome sandstone, Jarlemai siltstone formations and top of the seawater intrusion directly from inversion results, and thereby derive hydrogeological parameters such as the aquifer thickness. However, interpretation of the location of the seawater intrusion toe (Bear et al., 1999) which is the inland tip of the seawater intrusion is not so straightforward. An objective assessment of the location of the seawater intrusion toe, which would then allow an assessment of movement of the seawater intrusion is therefore of critical importance. Such an objective assessment is our main reason for employing a machine-learning algorithm. This is particularly important because once trained, subsequent surveys could be analysed using the same methodology in order to reduce the risk of interpreter bias.

RESULTS AND DISCUSSION

We evaluate algorithm performance by comparing results of classification to results of interpretation. We used the four lines indicated in Figure 1 in our evaluation. One of these lines (L 1006601) was chosen to form the training data, while the remaining three (L 1002501, L 1005001 and L 1008001) were used as the test data sets. Experiments showed little practical difference whether a line was taken as training or test data. We use confusion matrices (Stehman, 1997) as the primary metric for classifier accuracy. Confusion matrices plot correct and incorrect classification of each category of the test data compared to training data. If all the test data were accurately classified, then the confusion matrix would only have entries along the diagonal, and those entries would correspond to the number of data in each category. Misclassified test data appear as off-diagonal elements. A crude measure of algorithm performance is to express the ratio of the confusion matrix's trace to the number of data available for classification as a percentage.

Because they are used to interpret underlying geology, we use a feature vector based on the inverted models over each interpreted region. Since we interpret geology from inverted lines sections, our feature vector is Log10 of both layer conductivity and thickness together with ground elevation. Figure 4 presents confusion matrices for each classification algorithm. The best-performing algorithm was Random Forest with 100% of test data accurately classified. The next-best performing algorithm was Nearest Neighbours with 99.82% of test data accurately classified. Although the Nearest Neighbours algorithm accurately classified all QBJ and S samples, four QBS samples were misclassified as QBJ. The algorithm with the poorest performance was Artificial Neural Network with 97.56% of test data accurately classified. The Artificial Neural Network algorithm misclassified three QBJ samples as QBS, and 50 S samples as QBS. That all algorithms worked very well suggests that machine learning algorithms can be used to classify inversion results in terms of known geology. Training a machine learning algorithm on the measured data instead of the inversion results would mean that only for the training data an inversion would need to be performed. In the next section, therefore, we will investigate the use of alternative feature vectors.

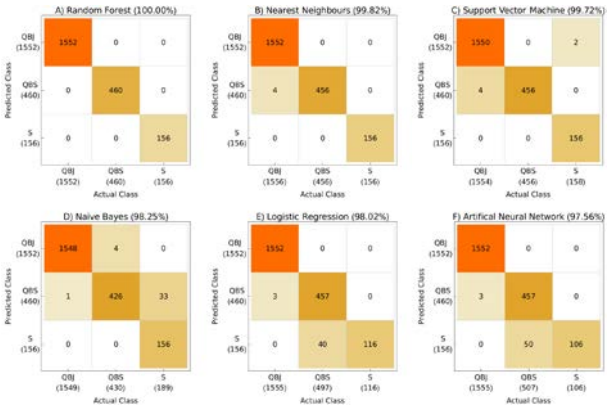

Figure 4: Comparison of results presented as confusion matrices.

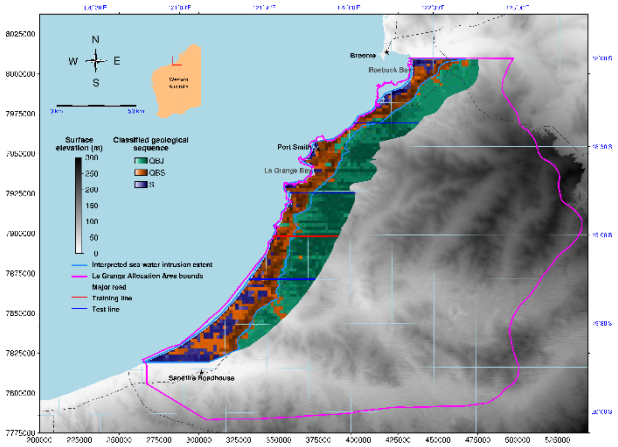

Figure 5: Results of the application of an RF algorithm trained to classify lithology in the LGGA. Comparison of the extent of the classified extent of SWI with that interpreted directly from layered-earth in version results is favourable.

Machine learning algorithms employed in this paper require numerical representation of features of interest. This numerical representation is termed the 'feature vector' and in the current context, the classifier derived from it represents an interpretation of geology. Because geology is interpreted from 1D layered-earth inversion results which are plotted as the variation of Log10 (layer conductivity) over depth, we encode interpreted geology using the inversion results over a range of stations associated with a known geology together with the ground elevation. Thus, our feature vector, based on a five-layer layered earth model, has 10 elements viz. logarithms of both layer conductivities and thicknesses and ground elevation. Ground elevation was included to aid classification; strong conductors near sea level are characteristic of seawater intrusion rather than (e.g.) clay layers which may appear at higher elevations. For the La Grange data, we found little practical difference when using either thicknesses or their logarithms in the feature vector.

Figure 5 presents results of applying a RF algorithm trained to classify lithology in the LGAA. Because the algorithm was trained using data near the coast, we restrict application to data from survey lines closest to the coast in order to minimise the chance of miss-classification of geology. Results are presented as probabilities results at a station being attributable to S, QBS or QBJ sequences. Salient points of Figure 5 are that the QBJ sequence is generally confined to the east of the extent of the seawater intrusion interpreted by Annetts et al. (2015). The S geological sequence is generally found close to the coast although there are some inland regions in the south of the LGAA associated with the Mandorah Marsh. Because the S geological sequence being defined using a near-surface high-conductivity zone, unless there is a significant contrast in conductivities over depth, we suggest that regions with thicker near-surface clays are likely to be (mis)classified as S. One way to ameliorate such misclassification might be to include the distance of a station from the coast in the feature vector. Tendrils of the QBS geological sequence extending into the OBJ (e.g. around 350000E between 7840000N and 7860000N) are associated with paleochannels. It is reasonable to speculate that the seawater intrusion front is more likely to move inland along the more permeable palaeochannels.

It is important to recognise that the feature vector may consist of other elements. For example, since model values are typically output by inversion codes are linearly-scaled, we might use them directly. Such a feature vector also has 10 elements. Nevertheless, conductivities are commonly plotted using logarithmic scales for qualitative interpretation and this motivates populating the feature vector with the logarithms of conductivities. Since they are commonly provided with data by contractors, we might use a feature vector consisting of conductivity values derived from a conductivity-depth transformation (CDT) (Macnae et al., 1998). Because layer thickness is constant for the CDT results provided with the La Grange survey, and the CDT used 5m thick layers over a 500m range of depths, a CDT-based feature vector consists of 101 elements. Since CDT's are often interpreted with Log10 (layer conductivity), we use Log10 (layer conductivity) in our CDT-based feature vector. The final alternative feature vector we consider consists of the secondary magnetic field response. Such a feature vector might be used in the event that the entire survey was not inverted or processed using a CDT. The TEMPEST survey over La Grange measured inline (the component of the secondary magnetic field response in the flight direction) and vertical components of the secondary magnetic field, therefore a feature vector consisting of data would consist of 31 elements (15 gates × two components + ground surface elevation). Clearly, the possibilities of feature-vector content are legion, indeed, feature vectors might combine different data sets. However, such variation was not considered here because of the excellent performance of the classifiers that was achieved with relatively-simple feature vectors.

Figure 6 presents results of a cross validation study. All machine-learning algorithms are compared using all feature vectors considered in this paper. Perfect classification is observed for a number of combinations of machine-learning algorithm and feature vector. Variation between individual techniques confirms the results of Caruana and Niculescu-Mizil (2006) that in practice, varying performance among machine-learning algorithms is to be expected. On average, using a feature vector consisting of Log10 (Model) allows accurate classification by all algorithms. Only an extremely

poor classification by the NB algorithm using a feature vector consisting of measured data prevents this from being recommended in general. We return to this point in Section 4. The central conclusion from Figure 6 is that only a small proportion of a large data set may need to be inverted and accurately interpreted in order to accurately classify geophysical data in terms of simple geology. Thus, inversion of the entire data set may not be required for its accurate classification in terms of known geology.

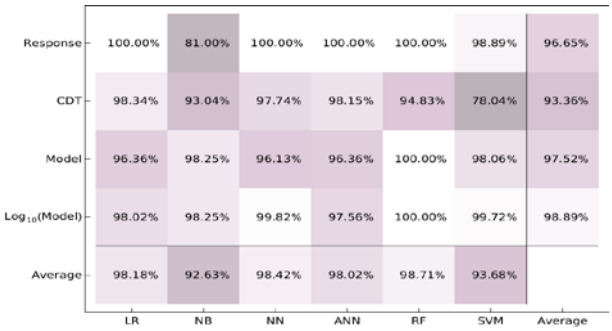

Figure 6: Comparison of ML algorithm accuracy with variation of feature vector. These results suggest that it may not be necessary to invert an entire data set in order to classify simple lithologies.

CONCLUSIONS

We have investigated the use of machine learning algorithms on geophysical data for geological interpretation. We investigated feature vector content and its effect on algorithm choice and accuracy. We found that after careful interpretation of a selection of test and training data sets, machine-learning algorithms could be successfully applied to geophysical data in the form of a number of different feature vectors. A major outcome of our study is the addition of a geological interpretation to results of geophysical inversions for a layered earth. The final result of our study is a probabilistic geological map of the La Grange which is a quantitative interpretation of a geophysical inversion results. A map of the inland extent of the seawater intrusion resulting from the probabilistic geological classification compared favourably with a manual line-by-line interpretation of inversion results. This lends support to the argument that machine learning techniques can at least speed up the interpretation of large AEM surveys if not automate it given a careful tuning and selection of algorithms. That all algorithms performed well is likely due to our careful interpretation of geophysical inversion results, and careful selection of a training data set. We showed that the results of such care, based on inversion results, could be used to accurately classify data that were not inverted. This may have applications where inversion is particularly expensive, for example, when multidimensional inversion is required, for very large data sets, or as a preliminary geological analysis. Because the feature vector is required to be accurately interpreted, relatively independent of its content, this offers the intriguing possibility for interpretation based on classification using a feature vector consisting of different, perhaps non-geophysical, data sets.

A major difference between a manual interpretation and the results of the machine learning algorithms is that they provide probabilities. When it comes to determining the extent of the sea water intrusion, maps of its probability can be used to identify areas where confidence into the assigned geological sequence is low or in areas of apparent miss-classification. Once such areas have been identified, a manual interpretation could be performed and be added to the training data and thus improve the overall predictive capabilities. Alternatively, if the collected data are inconclusive, then this would help identify areas for ground followup surveys.

ACKNOWLEDGEMENTS

This work would not have been possible without the groundwork laid by MRIWA project M435 (supported by Cameco Pty. Ltd., First Quantum Pty. Ltd., and Rio Tinto Exploration Pty. Ltd.). We thank the Department of Food and Agriculture, Western Australia (DAFWA) for making the field data available to the public. We also thank the Deep Earth Imaging Future Science Platform (DEI-FSP) for its support.

REFERENCES

Annetts, D., T. Munday, R. George, N. Wright, and K. Cahill, 2016, The application of AEM to mapping the aquifer and groundwater characteristics of the La Grange allocation area, WA: Final report: Technical Report EP165299, CSIRO, Kensington, WA.

Bear, J., A.-D. Cheng, S. Sorek, D. Ouazar, and I. Herrera, eds., 1999, Seawater intrusion in coastal aquifers – concepts, methods and practices: Springer-460 Science+Business Media B.V.

Bishop, C., 1995, Neural networks for pattern recognition: Clarendon Press.

Breiman, L., 2001, Random Forests: Machine Learning, 45, 5–32.

Brodie, R. C., 2016, GALEISBSTDEM: A deterministic algorithm for 1d sample by sample inversion of time-domain AEM data: theoretical details, Geoscience Australia, https://github.com/geoscienceaustralia/ga-aem. Accessed 23 November, 2016.

Caruana, R., and A. Niculescu-Mizil, 2006, An Empirical Comparison of Supervised Learning Algorithms Using Different Performance Metrics: Presented at the 23rd international conference on Machine learning, ACM.

Cortes, C., and V. Vapnik, 1995, Support-vector networks: Machine Learning, 20, 273–297.

Cover, T., and P. Hart, 1967, Nearest neighbor Pattern classification: IEEE Transactions on Information Theory, IT-13, 21–27.

Cox, D. R., 1958, The Regression Analysis of Binary Sequences: Journal of the Royal Statistical Society. Series B (Methodological), 20, 215–242.

Cracknell, M., and A. 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.

Dumay, J., and F. Fournier, 1988, Multivariate statistical analyses applied to seismic facies recognition: Geophysics, 53, 1151–1159.

Forman, D., and D. Wales, 1981, Geological evolution of the Canning Basin, 483 Western Australia: Bulletin 210, Bureau of Mineral Resources, Geology and Geophysics, Canberra, Australia.

Ghassemi, F., H. Etminan, and J. Ferguson, 1992, A reconnaissance investigation of the major Palaeozoic aquifers in the Canning Basin, Western Australia, in relation to Zn-Pb mineralisation: BMR Journal of Australian Geology and Geophysics, 13, 37–57.

Hall, P., B. Park, and R. Samworth, 2008, Choice of neighbour order in nearest neighbour classification: The Annals of Statistics, 36, 2135–2152.

Ho, T., 1995, Random decision forests: Presented at the Third International Conference on Document Analysis and Recognition, IEEE.

Lane, R., A. Green, C. Golding, M. Owers, P. Pik, C. Plunkett, D. Sattel, and B. Thorn, 2000, An example of 3D conductivity mapping using the TEMPEST airborne electromagnetic system: Exploration Geophysics, 31, 162–172.

Lau, J., D. Commander, and G. Jacobson, 1987, Hydrogeology of Australia: Bulletin 227, Bureau of Mineral Resources, Geology and Geophysics, Canberra, Australia.

Lin, Y., and Y. Jeon, 2006, Random forests and adaptive nearest neighbors: Journal of the American Statistical Association, 101, 578–590.

Macnae, J., A. King, N. Stolz, A. Osmakoff, and A. Blaha, 1998, Fast AEM data processing and inversion: Exploration Geophysics, 29, 163–169.

Marker, P. A., N. Foged, X. He, A. V. Christiansen, J. C. Refsgaard, E. Auken, and P. Bauer-Gottwein, 2015, Performance evaluation of groundwater model hydrostratigraphy from airborne electromagnetic data and lithological borehole logs: Hydrology and Earth System Sciences, 19, 3875–3890.

MIRA Geoscience, 2015, Inversion and preliminary interpretation of airborne electromagnetic data from the West Canning Basin Sandfire TEMPEST survey: Technical report, MIRA Geoscience, Perth, Western Australia.

Ng, A., and Jordan, 2001, On Discriminative vs. Generative classifiers: A comparison of logistic regression and naive Bayes: Presented at the Advances in Neural Information Processing Systems 14.

Niculescu-Mizil, A., and R. Caruana, 2005, Predicting good probabilities with supervised learning: Presented at the 22nd international Conference on Machine Learning.

Paul, R., R. George, and P. Gardiner, 2013, A review of the Broome sandstone aquifer in the La Grange area: Resource Management Technical Report 387, Department of Agriculture and Food, Perth, Western Australia.

Prinzie, A., and D. Van den Poel, 2008, Random Forests for multiclass classification: Random MultiNomial Logit: Expert Systems with Applications, 34, 1721–1732.

Raiche, A., 1991, A pattern recognition approach to geophysical inversion using neural nets: Geophys. J. Int, 105, 629–648.

Saggaf, M. M., M. N. Toksöz, and M. I. Marhoon, 2003, Seismic facies classification and identification by competitive neural networks: Geophysics, 68, 1984–1999.

Samworth, R., 2012, Optimal weighted nearest neighbor classification: The Annals of Statistics, 40, 2733–2763.

Stehman, S. V., 1997, Selecting and interpreting measures of thematic classification accuracy: Remote Sensing of Environment, 62, 77–89.

Vapnik, V., 1982, Estimation of dependencies based on empirical data, Addendum 1: Springer-Verlag.

Wolfram Research Inc., 2016, Mathematica.

Zen, H., and H. Sak, 2015, Unidirectional Long Short-Term Memory Recurrent Neural Network with Recurrent Output Layer for Low-Latency Speech Synthesis: Presented at the Proceedings of International Conference on Accoustics, Speech and Signal Processing, IEEE.

Zhang, H., 2004, The Optimality of Naive Bayes: Presented at the Seventeenth International Florida Artificial Intelligence Research Society Conference.

Zuo, R., and E. J. M. Carranza, 2011, Support vector machine: A tool for mapping mineral prospectivity: Computers and Geosciences, 37, 1967–1975.