# Interpretation of Geophysical Surveys of Archaeological Sites using Artificial Neural Networks

David J. Bescoby
Institute of World Archaeology,
University of East Anglia,
Norwich, U.K.
d.bescoby@uea.ac.uk

Gavin C. Cawley
School of Information Systems,
University of East Anglia,
Norwich, U.K.
gcc@sys.uea.ac.uk

P. Neil Chroston
School of Environmental Sciences,
University of East Anglia,
Norwich, U.K.
p.chroston@uea.ac.uk

***Abstract*— Geophysical surveys form a vital component of modern archaeological studies, providing data on the location of surviving sub-surface remains. Unfortunately interpretation of the raw output from geophysical surveys requires considerable expertise and current methods for machine interpretation are computationally expensive. In this paper, we describe the use of multi-layer perceptron neural networks for automated interpretation of geophysical survey data and present promising initial results for a survey taken of the Roman suburb of Butrint, a classical city situated in what is now Albania. Clearly a computationally efficient and straightforward means of inverting geophysical survey data is potentially an extremely useful tool in encouraging public awareness of cultural heritage.**

## I. Introduction

The classical city of Butrint, located upon the eastern shore of the Straits of Corfu in Southern Albania, encompasses over 3000 years of Mediterranean history and archaeology. Soon after 31 BC, following victory at the battle of Actium, the emperor Augustus declared the city of Butrint a colony. The archaeological record shows extensive construction in and around the city at this period which has become the focus of the current geophysical research. The city of Butrint is located upon a small limestone promontory jutting out into a lagoonal lake, the remnants of a former coastal embayment. The city is today flanked to the south and west by a large deltaic alluvial plain and is now over 1 km inland, connected to the coast by a narrow channel draining from the lake. The colonial expansion of the city focused strongly upon the developing floodplain, adding significantly to its original size. Continued accretion has led to the burial of this phase of settlement since its abandonment sometime after the 6th century AD. In order to investigate the nature and extent of the former Roman suburb, an extensive archaeological geophysical survey has been undertaken in order to map surviving sub-surface remains. A high resolution total-field magnetic gradiometer survey (see fig 1), measuring small fluctuations in the earths' magnetic field caused by subsurface anomalies, has revealed the buried remains of an extensive planned settlement and a number of outlying buildings beneath the floodplain.

The majority of archaeological geophysical surveys are concerned solely with identifying the location and extent of buried archaeological remains. However, much geophysical data inherently contains information relating to the depth and geometry of anomaly features which is currently ignored due to the lack of tools for automated interpretation. The archaeological site at Butrint was inscribed as a World Heritage Site in 1992, and the area enlarged in 2000 to encompass much of the Vrina Plain to the south, upon which the current study is focused. As a protected cultural heritage resource, potential for the archaeological excavation of buried remains is small. As a result archaeological investigations necessarily rely on geophysical remote sensing techniques to map the location and extent of surviving sub-surface remains. It is therefore vital to extract the maximum amount of information from the geophysical data, in terms of clarifying the location, extent and burial depth of features of interest. This not only has intrinsic archaeological value, but is also vitally important to the ongoing management strategies of the cultural landscape in terms of current and future land-use.

Fig. 1. One of the authors (DJB) performing the geomagnetic survey of the Roman suburb at the Butrint site in Southern Albania. Measurements were made at 0.25m intervals along lines 1m apart.

The aim of the current research is to model the depth and shape of source anomalies from the magnetic measurements, providing an enhanced interpretation of the data through a sub-surface reconstruction of archaeological features. While it is relatively straight-forward to generate simulated geophysical survey data for a given set of buried remains, taking into account a model of the magnetic properties of the soil covering the site, “inversion” of geomagnetic data to discover the depth of sub-surface features has proved considerably more difficult. The inversion of magnetic data has therefore been traditionally achieved by constructing a magnetic sub-surface model of the predicted features, composed of a regular array of homogeneous dipole sources or blocks having a specified value of magnetisation [1, 2]. The determination of a final sub-surface model for a given set of magnetic measurements is then achieved by adjusting the magnetisation of the dipole sources until a satisfactory approximation of the measured data is achieved [3]. The reconstruction problem is therefore one of optimisation which can be solved by a wide variety of iterative optimisation algorithms and global search methods [4, 5]. Plausible models are achieved by imposing a number of *a priori* constraints often derived from ground truth data. Clearly this is a computationally expensive procedure. Incomplete and noisy data often cause problems in achieving a consistent solution, especially when resolving complicated archaeological structures.

The work presented here investigates a fundamentally different approach to deriving a sub-surface model from magnetic survey data, by utilising a simple multi-layer perceptron neural network to learn the non-linear mapping between measured data and sub-surface model parameters. Artificial neural networks have been successfully applied to a number of other geophysical modelling problems, including parameter prediction and estimation, classification, filtering and optimisation [6–9], however neural networks have not yet been widely applied in computational archaeology. The adoption of a machine learning approach potentially holds several advantages over more conventional modelling techniques, allowing the effective solution of non-linear problems with complex, incomplete and noisy data. The computational expense for neural inversion, being dependent upon the dimension of the space of unknown parameters, rather than the physical dimensions of medium, is very low [10].

The remainder of this paper is structured as follows: In section II we describe the neural network architecture used, including details of the Bayesian regularisation scheme used to avoid over-fitting. Section III describes the construction of a neural network for inversion of geophysical survey data, using simulated magnetic survey data. Results obtained by inversion of real magnetic survey data for the Roman suburb of Butrint are presented in section IV.

## II. Neural Network Architecture

For this study, we adopt the familiar Multi-Layer Perceptron network architecture (see e.g. Bishop [11]). The optimal model parameters, $\boldsymbol{w}$, are determined by gradient descent optimisation of an appropriate error function, $E_\mathcal{D}$, over a set of training examples, $\mathcal{D} = \{(\boldsymbol{x}_i, t_i)\}_{i=1}^N$, $\boldsymbol{x}_i \in \mathcal{X} \subset \mathbb{R}^d$, $t_i \in \mathbb{R}$, where $\boldsymbol{x}_i$ is the vector of explanatory variables and $t_i$ is the desired output for the $i^{th}$ training pattern. The error metric most commonly encountered in non-linear regression is the sum-of-squares error, given by

$$E_\mathcal{D} = \frac{1}{2}\sum_{i=1}^N (y_i - t_i)^2, \tag{1}$$

where $y_i$ is the output of the network for the $i^{th}$ training pattern. In order to avoid over-fitting to the training data, however, it is common to adopt a regularised [12] error function, adding a term $E_\mathcal{W}$ penalising overly-complex models, i.e.

$$M = \alpha E_\mathcal{W} + \beta E_\mathcal{D}, \tag{2}$$

where $\alpha$ and $\beta$ are regularisation parameters controlling the bias-variance trade-off [13]. Minimising a regularised error function of this nature is equivalent to the Bayesian approach which seeks to maximise the posterior density of the weights (e.g. [14, 15]), given by

$$P(\boldsymbol{w} \mid \mathcal{D}) \propto P(\mathcal{D} \mid \boldsymbol{w})P(\boldsymbol{w}),$$

where $P(\mathcal{D} \mid \boldsymbol{w})$ is the likelihood of the data and $P(\boldsymbol{w})$ is a prior distribution over $\boldsymbol{w}$. The form of the functions $E_\mathcal{D}$ and $E_\mathcal{W}$ correspond to distributional assumptions regarding the data likelihood and prior distribution over network parameters respectively. The usual sum-of-squares metric (1), corresponds to a Gaussian likelihood,

$$P(\mathcal{D} \mid \boldsymbol{w}) = \frac{1}{\sqrt{2\pi\beta^{-1}}} \exp\left\{-\frac{[t_i - y(\boldsymbol{x}_i)]^2}{2\beta^{-1}}\right\}$$

with fixed variance $\sigma^2 = 1/\beta$. For this study, we adopt the Laplace prior propounded by Williams [16], which corresponds to a $L_1$ norm regularisation term,

$$E_\mathcal{W} = \sum_{i=1}^W |w_i|. \quad \iff \quad P(w) = \frac{1}{2\beta}\exp\left\{-\frac{|w|}{\beta}\right\}$$

where $W$ is the number of model parameters. An interesting feature of the Laplace regulariser is that it leads to pruning of redundant model parameters. From 2, at a minimum of $M$ we have

$$\left|\frac{\partial E_y}{\partial w_i}\right| = \frac{\alpha}{\beta} \qquad w_i > 0, \qquad \left|\frac{\partial E_y}{\partial w_i}\right| < \frac{\alpha}{\beta} \qquad w_i = 0.$$

As a result, any weight not obtaining the data misfit sensitivity of $\alpha/\beta$ is set exactly to zero and can be pruned from the network.

### A. Eliminating Regularisation Parameters

The hyperparameters $\alpha$ and $\beta$ can be estimated by maximising the evidence [14] or alternatively may be integrated out analytically [16, 17]. Here we take the latter approach; the posterior distribution of the parameters is given by

$$p(\boldsymbol{w}) = \int p(\boldsymbol{w}|\alpha)p(\alpha)d\alpha. \tag{3}$$

Assuming the Laplace prior, the prior distribution over the weights of the network, conditioned on the regularisation parameter $\alpha$, is given by,

$$p(\boldsymbol{w}|\alpha) = Z_{\mathcal{W}}(\alpha)^{-1} \exp\{-\alpha E_{\mathcal{W}}\} \tag{4}$$

where the necessary normalising constant is given by

$$Z_{\mathcal{W}}(\alpha) = \left(\frac{2}{\alpha}\right)^{W}. \tag{5}$$

Substituting equations 4 and 5 into equation 3, adopting the (improper) uninformative Jeffreys prior, $p(\alpha) = 1/\alpha$ [18], and noting that $\alpha$ is strictly positive,

$$= \int_0^\infty 2^{-W} \alpha^{W-1} \exp\{-\alpha E_{\mathcal{W}}\} d\alpha.$$

Using the Gamma integral, $\int_0^\infty x^{\nu-1} e^{-\mu x} dx = \frac{\Gamma(\nu)}{\mu^\nu}$ (Gradshteyn and Ryzhik [19], equation 3.384), we obtain

$$p(\boldsymbol{w}) = \frac{\Gamma(W)}{(2E_{\mathcal{W}})^W}.$$

Taking the negative logarithm and omitting irrelevant constant terms,

$$-\log p(\boldsymbol{w}) = W \log E_{\mathcal{W}}. \tag{6}$$

Applying a similar treatment to the data misfit term (assuming a sum-of-squares error), we have

$$L = \frac{1}{2} N \log E_{\mathcal{D}} + W \log E_{\mathcal{W}}.$$

For a network with more than one output unit, it is sensible to assume that each output has a different noise process (and therefore a different optimal value for $\beta$). It is also sensible to assign hidden layer weights and weights associated with each output unit to different regularisation classes so they are regularised separately. This leads to the training criterion used in this study:

$$L = \frac{N}{2} \sum_{i=1}^{O} \log E_{\mathcal{D}}^i + \sum_{j=1}^{C} W_j \log E_{\mathcal{W}}^j,$$

where $O$ is the number of output units, $C$ is the number of regularisation classes (groups of weights sharing the same regularisation parameter) and $W_j$ is the number of non-zero weights in the $j^{th}$ class. Note that bias parameters are not normally regularised.

### B. Choice of Data Misfit Term

While the conventional sum-of-squares misfit term would be appropriate for this study, we adopt a data misfit term corresponding to a heteroscedastic (input dependent variance) Gaussian noise process, i.e.

$$E_{\mathcal{D}} = \sum_{i=1}^{\ell} \left\{ \log \sigma(\boldsymbol{x}_i) + \frac{[\mu(\boldsymbol{x}_i) - t_i]^2}{2\sigma^2(\boldsymbol{x}_i)} \right\}. \tag{7}$$

Note the multi-layer perceptron network now has two output units, one giving the mean of the target distribution, $\mu(\boldsymbol{x})$, as before, and an additional unit giving the predicted variance, $\sigma(\boldsymbol{x})$. A linear activation function is used in the output unit corresponding to $\mu(\boldsymbol{x})$, and an exponential activation function for the unit corresponding to $\sigma(\boldsymbol{x})$, to enforce strictly positive estimates of conditional variance. This approach provides two advantages: Firstly the estimates of conditional variance provide error bars, indicating the uncertainty of model predictions [20–22]. Secondly the output of the model now completely specifies the target distribution, so the regularisation parameter $\beta$ is no longer necessary.

## III. Method

Suitable training data were derived from a range of synthetic models of buried wall foundations of a variety of types and burial depths and the corresponding magnetic responses produced via forward modelling. The synthetic models were based upon known archaeological examples recorded in a number of trial excavations conducted within the study area, in which surviving wall foundations were found to be the dominant feature of interest. The forward modelling of magnetic responses corresponding to each model was undertaken using software developed by the University of British Columbia under a consortium research project [23]. The measurable distribution of magnetisation $\boldsymbol{J}$ caused by a buried anomalous feature, such as a wall foundation, is proportional to the magnetic susceptibility of the feature, $\kappa$, within the inducing magnetic field of the earth, $\boldsymbol{H}_o$, i.e.

$$\boldsymbol{J} = \kappa \boldsymbol{H}_o, \tag{8}$$

see [23]. The resulting anomalous magnetic field as a result of the magnetism of the feature, $\boldsymbol{J}$, is calculated by the following integral equation, see [23]

$$\boldsymbol{B}_a(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_V \nabla\nabla \frac{1}{\boldsymbol{r} - \boldsymbol{r}^l} . \boldsymbol{J} \delta v \tag{9}$$

where $\boldsymbol{r}$ is the location of the observation, $\mu_o$ is the permeability of free space and $V$ is the volume of the magnetising feature. Each model is generated within a 3D orthogonal mesh of cuboidal cells representing the earths sub-surface, by defining the magnetic susceptibility of each cell. The resulting anomalous magnetic field is then measured at a series of equally spaced observation points at a height above the modelled surface, the resulting magnetic field being the sum of fields produced by all cells having a non-zero susceptibility value [23]. Appropriate values of the magnetic susceptibility of modelled walls and associated deposits were derived from laboratory measurements of representative building materials.

The modelled data was input into the network as a vector of magnetic field values, representing an effective 2 metre by 2 metre area above the subsurface model. The corresponding target data consisted of a single value representing the depth below the surface to the modelled wall features, falling at the centre of the input magnetic field values. Subsequent training pairs were derived by moving the 2 x 2 metre input window sequentially over the modelled area. Noise sampled from the magnetic survey data was then added to the training data to create a more realistic response, following the method outlined by Schollar (1970) [24]. The Fourier transform of sampled magnetic noise is randomised by computing the modulus and the angle of the transform, and adding a randomly distributed number to the angle, and recalculating the Fourier coefficient. When the inverse transform is taken, the resulting noise anomalies retain the power spectrum of the original, but in differing distributions [24].

Some pre-processing of the input training data as also undertaken to enhance spatial symmetry of the magnetic data and introduce rotational invariance. A reduction to the Pole operator was applied to the magnetic data to improve the overall symmetry of magnetic anomalies. This operation transforms the magnetic response caused by an arbitrary source into the magnetic anomaly that the same source would produce if it were located at the Earth's (north) pole, where the inducing magnetic field is vertical. Rotational invariance was introduced in order to reduce the amount of training data required to represent all possible wall orientations. This was achieved by computing the mean average of a series of concentric rings centred upon each 2m by 2m square of input data, effectively deriving a rotationally invariant input representation. In addition a range of training data representing unwanted magnetic responses such as those caused by modern land drains were added to the training data set, with corresponding null target data.

## IV. Results

The synthetic data used in training and testing the neural network geophysical interpretation system is based on an excavated building from the Butrint site. The synthetic remains are shown in figure 2 (a), along with geomagnetic survey data simulated via forward-modelling, as described in the previous section (b) and predicted subsurface features (c). The root-mean-square error for the test set is 0.25m.

Figure 3 shows the geomagnetic survey of an outlying single phase Roman building (i.e. it represents a single phase of cultural activity) (a), along with a shaded relief plot of predicted subsurface remains (b) and picture of a trial trench (c) excavated within the same area. The building is located approximately 100m to the north east of the main suburb shown in figure 4. The shaded relief plot gives strong indication of the structure of the surviving wall foundations of the building, partially verified by the excavation of a small trench, providing some "ground truth" data. The RMS error for the predicted depth of buried remains in this trench is 0.23m, which is an acceptable level of accuracy for archaeological purposes (figure 3 (c) shows the remains following the removal of a comparable amount of loose material).

Figure 4 shows the raw geomagnetic survey data (a) and shaded relief plot of predicted subsurface remains (b) overlaid onto a satellite image of the study area. Clearly the shaded relief plot gives a far more vivid impression of the buried structure than the raw geomagnetic data. This type of result is likely to be very helpful in encouraging public understanding of cultural heritage sites. The model can be seen to predict significantly more shallow burial depths for remains centred around the north western portion of the survey, a result which was later confirmed by a series of boreholes. Although topographically flat today, this area is formed upon a raised gravel bank, and was once a natural topographical high. The shallower burial depths over this area seem to relate to the thinner cover of alluvial deposits. It is also possible that the shallower wall depths reflect prolonged occupation of the higher ground, less prone to flooding events during periods of increased sea level, known to have occurred during the 5th century AD.

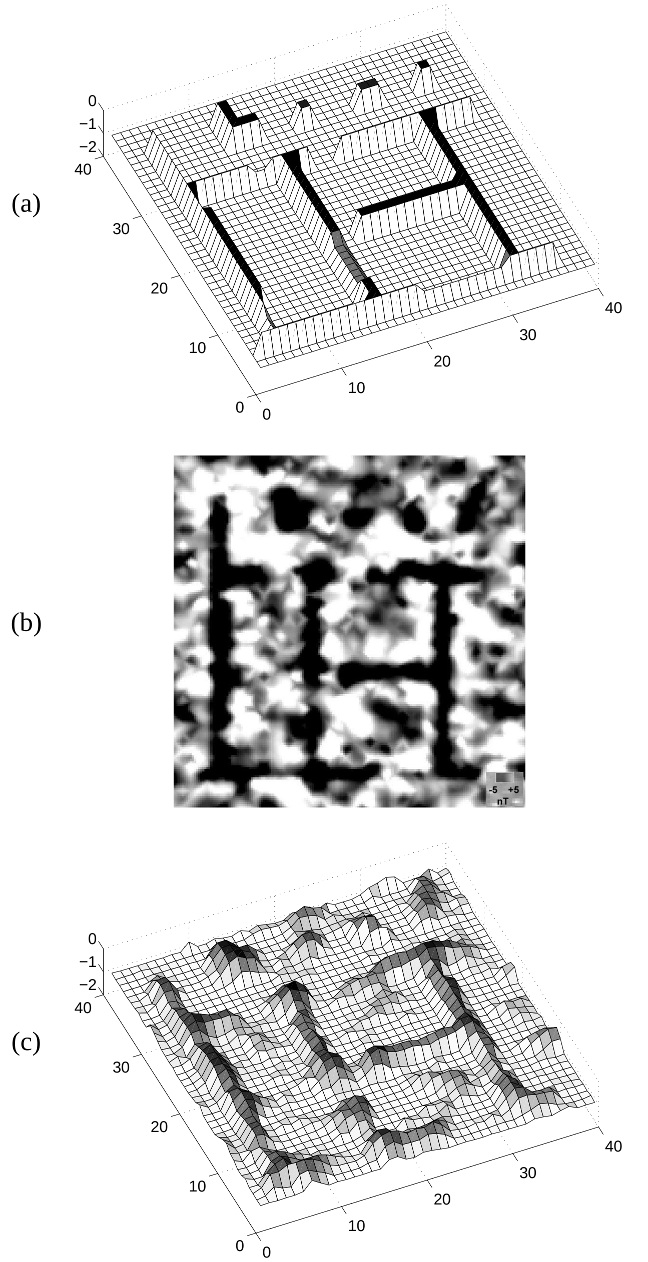


Fig. 2. Synthetic test data based on excavated building from the Butrint site (a), example simulated geomagnetic survey data (b) and subsurface features predicted by neural network model (dimensions in metres).

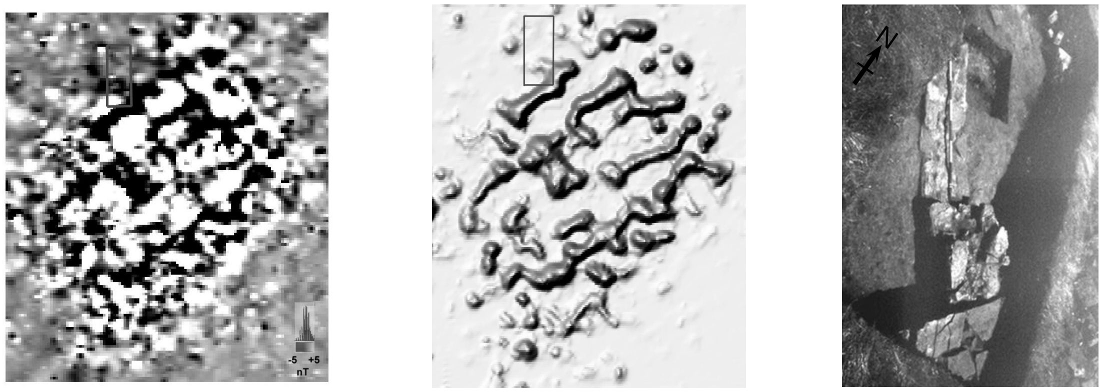


(a) (b) (c)

Fig. 3. Geomagnetic survey of outlying single phase Roman building (a), shaded relief plot of predicted subsurface remains (b) and $5 \times 2$ metre trial trench excavated (c) corresponding to area marked on (a) and (b).


(a) (b)

Fig. 4. Raw geomagnetic survey data (a) overlayed onto a pan-sharpened satellite image (copyright DigiGlobe 2002) of the Butrint site and (b) shaded relief plot of the depth predictions from neural network committee.

## V. Conclusion

Geophysical surveys have become a integral component of modern archaeology, however the resulting data is rarely fully exploited, often being used only to give and indication of the location and extent of buried remains. In this study we have investigated the use of artificial neural networks for further interpretation of geomagnetic survey data to give an indication of the depth of surviving wall foundations. Initial results for a survey of a Roman suburb of the classical city of Butrint appear highly promising. The authors hope that this work will prove particularly useful in encouraging public interest in the archaeology of landscapes, which would otherwise appear featureless, such as that shown in figure 1. This is especially important for world heritage sites such as Butrint where large-scale excavation is not permitted.

## Acknowledgements

The authors would like to thank Graham Finlayson for improving the quality of the image seen in figure 1, Nicola Talbot for her assistance in typesetting this paper and Mike Lincoln and Andy Hanna for their helpful comments on previous drafts. This work was supported by the Natural Environment Research Council (grant number NER/S/A/2000/03321A) and by the Butrint Foundation (registered charity no. 1017039).

## References

[1] Y. Li and D. W. Oldenburg, "3-D inversion of magnetic data," *Geophysics*, vol. 61, no. 2, pp. 394–408, Mar.-Apr. 1996.

[2] I. Schollar, "A program for the simulation of magnetic anomalies of archaeological origin in a computer," *Prospezioni Archeologiche*, vol. 4, pp. 59–83, 1969.

[3] M. Doneus, A. Eder-Hinterleitner, and W. Neubauer, "Archaeological prospection in Austria," in *Archaeological Prospection*, M. Doneus, A. Eder-Hinterleitner, and W. Neubauer, Eds., Wien, Austria, Sept.19–23 2001, Fourth International Conference on Archaeological Prospection, pp. 11–33, Austrian Academy of Sciences.

[4] J. Herwanger, H. Maurer, A. G. Green, and J. Leckebusch, "3-D inversion of magnetic gradiometer data in archaeological prospecting: Possibilities and limitations," *Geophysics*, vol. 65, no. 3, pp. 849–860, May-June 2000.

[5] A. Eder-Hinterleitner, W. Neubauer, and P. Melichar, "Reconstruction of archaeological structures using magnetic prospection," *Analecta praehistorica Leidensia*, vol. 28, pp. 131–137, 1996.

[6] M. M. Poulton, B. K. Sternberg, and C. E. Glass, "Location of subsurface targets in geophysical data using neural networks," *Geophysics*, vol. 57, no. 12, pp. 1534–1544, Dec. 1992.

[7] C. Calderón-Macías, M. K. Sen, and P. L. Stoffa, "Artificial neural networks for parameter estimation in geophysics," *Geophysical Prospecting*, vol. 48, pp. 21–47, 2000.

[8] M. van der Baan and C. Jutten, "Neural networks in geophysical applications," *Geophysics*, vol. 65, no. 4, pp. 1032–1047, July-Aug. 2000.

[9] G. El-Qady and K. Ushijima, "Inversion of DC resistivity data using neural networks," vol. 49, pp. 417–430, 2001.

[10] V. Spichak and I. Popova, "Artificial neural network inversion of magnetotelluric data in terms of three-dimensional earth macroparameters," *Geophysical Journal International*, vol. 142, pp. 15–26, 2000.

[11] C. M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.

[12] A. N. Tikhonov and V. Y. Arsenin, *Solutions of ill-posed problems*, John Wiley, New York, 1977.

[13] S. Geman, E. Bienenstock, and R. Doursat, "Neural networks and the bias/variance dilemma," *Neural Computation*, vol. 4, no. 1, pp. 1–58, 1992.

[14] D. J. C. Mackay, "A practical Bayesian framework for backprop networks," *Neural Computation*, vol. 4, pp. 448–472, 1992.

[15] R. M. Neal, *Bayesian Learning for Neural Networks*, Springer, New York, 1996.

[16] P. M. Williams, "Bayesian regularisation and pruning using a Laplace prior," *Neural Computation*, vol. 7, no. 1, pp. 117–143, 1995.

[17] W. L. Buntine and A. S. Weigend, "Bayesian back-propagation," *Complex Systems*, vol. 5, pp. 603–643, 1991.

[18] H. S. Jeffreys, *Theory of Probability*, Oxford University Press, 1939.

[19] I. S. Gradshteyn and I. M. Ryzhic, *Table of Integrals, Series and Products*, Academic Press, fifth edition, 1994.

[20] D. A. Nix and A. S. Weigend, "Estimating the mean and variance of the target probability distribution," in *Proc., Int. Conf. on Neural Networks*, 1994, vol. 1, pp. 55–60.

[21] D. A. Nix and A. S. Weigend, "Learning local error bars for nonlinear regression," in *Advances in Neural Information Processing Systems*, vol. 7, pp. 489–496. MIT Press, 1995.

[22] P. M. Williams, "Using neural networks to model conditional multivariate densities," *Neural Computation*, vol. 8, pp. 843–854, 1996.

[23] UBC-Geophysical Inversion Facility, *A program library for forward modelling and inversion of magnetic data over 3D structures*, UBC-Geophysical Inversion Facility, Department of Geophysics and Astronomy, University of British Columbia, Vancouver, British Columbia, 1998.

[24] I. Schollar, "Fourier transform methods for the evaluation of magnetic maps," *Prospezioni Archeologiche*, vol. 5, pp. 9–41, 1970.