

Neural network reconstruction for tomography of a gravel–air–seawater mixture

G Teague¹, J Tapson¹ and Q Smit²

¹ Department of Electrical Engineering, University of Cape Town, UCT Private Bag, Rondebosch 7701, South Africa

² Electrical Engineering, Cape Technikon, PO Box 652, Zonnebloem 8000, South Africa

Received 14 February 2001, in final form 18 April 2001, accepted for publication 2 May 2001

Abstract

This paper presents the first practical implementation of electrical impedance tomography for the imaging of a gravel–air–seawater mixture, wherein the image reconstruction is performed using a neural network. Although research into the use of neural networks for image reconstruction has been done, previous work has predominantly made use of simulated capacitance or impedance readings to train the networks. In contrast, this paper adopts a new practical approach whereby the networks are trained and tested using real data from an existing impedance tomography system. The practical aspects of generating this database are discussed and the network training details are given. Results indicate that the trained networks are able to discriminate among the seawater, gravel and air phases and reconstruct images of bubble configurations not included in the training database. In terms of the intended application, accurate predictions of volume fraction are required. Consequently, this paper examines an alternative solution to this problem, whereby the networks are trained to predict the volume fractions directly rather than first performing an image reconstruction as is standard. Results indicate that these volume fraction predictors outperform the image reconstruction networks in terms of the accuracies of the volume fraction predictions for each of the three phases.

Keywords: impedance tomography, neural network, image reconstruction

1. Introduction

Tomography can be described as the measurement of some characteristic by examining it in cross-section [1]. A tomographic reconstruction of the cross-section of a vessel can be achieved using a variety of measurement variables, including resistance, capacitance and impedance. In particular, capacitance tomography is a technique used to reconstruct the dielectric distribution within a vessel non-invasively. It operates on the principle that the capacitance of a pair of electrodes depends on the dielectric distribution of the material between the electrodes. By mounting several of these electrodes on the periphery of the vessel, it is possible to reconstruct an image of the contents of the vessel. Resistance tomography constructs an image of the conductivity distribution within the cross-section of the

vessel [2] and is intended for the imaging of multi-phase conductive components [3].

It is the task of the reconstruction algorithm to determine an image of the contents of the vessel from these limited measurements [4]. The standard reconstruction technique is linear back-projection. However, this technique is flawed and errors in the reconstruction occur due to the ‘soft-field’ effect [5]. This paper will examine the use of a neural network as an alternative image reconstruction technique.

Research into the use of neural networks for image reconstruction in capacitance tomography has been done [4–7]. However, previous work has predominantly made use of simulated systems for which the capacitance readings are obtained using a finite element modelling technique. The same is also true for impedance tomography systems [8, 9]. In contrast, this paper will examine the use of neural networks

from a practical standpoint in that the networks will be trained and tested using real data from a physical capacitance tomography system.

The objective of this work is to develop an instrument to measure the volume fractions of gravel, air and seawater in a pipeline. A neural network will be used to estimate these fractions from the capacitance and conductance readings obtained from the impedance tomography system.

The results of this paper show that a neural network is capable of performing a reconstruction of a gravel–air–seawater mixture using readings from a physical tomography system. The practical aspects of this task will be examined. In addition, a simple modification to the neural network architecture which improves the accuracy of the volume fraction predictions will be discussed.

2. The tomographic system hardware

A rig was constructed to simulate a static situation within a pipeline. It consists of a polyester pipe of inner diameter 0.22 m and height 0.3 m. One end of the pipeline is sealed so that it retains seawater when the axis of the pipeline is vertical. The electrode array consists of eight plate electrodes mounted on the external periphery of the vessel. These electrodes are connected to the data acquisition system, which performs the task of switching between the 28 linearly independent electrode combinations and measuring the corresponding capacitances. Synchronous detection circuitry is used to produce an output voltage proportional to the capacitance between the driver and receiver electrodes [10,11]. A microcontroller controls the switching task and performs an analogue-to-digital conversion on the resulting output voltages. These 10-bit values are then transmitted to the reconstruction computer via a serial link [12].

Initial tests conducted using this system produced promising results. However, it is not ideal to use capacitance tomography for the imaging of conductive substances. In addition, Stott *et al* [13] concluded that externally mounted electrodes introduce additional complications. It was shown that externally mounted electrodes produce satisfactory results only for substances of both low conductivity and low permittivity. However, the predominant phase of the current system is seawater, which is both conductive and has a relatively high dielectric constant. It was therefore decided to extend the system to include resistance tomography and, in so doing, develop an impedance tomography system.

The standard measuring technique for resistance tomography is the ‘four-electrode adjacent pair measurement protocol’ and operates by injecting an alternating current from one electrode to another for an adjacent pair of electrodes. The voltages between all adjacent pairs of electrode are then measured for the remaining electrodes [3]. The current injection is then moved to the next pair of electrodes and the process is repeated. The motivation for the adjacent measurement protocol is to ensure that any variation in the contact impedance between the electrodes and the substance being imaged does not influence the measured voltage [3]. However, if the contact impedance is negligible relative to the resistivity of the substance being imaged, then the standard two-electrode technique can be employed [3].

Stainless steel point electrodes were installed at the centres of the original plate electrodes. Since the contact impedance between these point electrodes and the seawater is negligible relative to the resistivity of the seawater, it was decided that the two-electrode technique could be used. In addition, the original capacitance tomography system employed the two-electrode technique. Consequently, the data acquisition system could be employed with little modification. The modifications required included extending the synchronous detection circuitry to measure the conductance between the driver and receiver electrode and programming the microcontroller to transmit an additional 28 conductance readings to the reconstruction computer. A reconstruction of the vessel contents had therefore to be achieved using these 28 capacitance and 28 conductance readings obtained from the data acquisition system.

3. Generating a training database

A neural network is a highly parameterized mathematical function that offers a general framework for representing non-linear functional mappings from several input variables to several output variables [14]. To estimate the volume fractions of the various phases within the pipeline, a pixel image of the cross-section of the vessel is first constructed. Each pixel in this image can be considered as the output of an individual neural network performing a classification function for that particular pixel, namely

$$\text{pixel} = f_{\text{CLASSIFY}}(\vec{C}, \vec{R}) \quad (1)$$

where

$$\text{pixel} \in \{\text{seawater, gravel, air}\}. \quad (2)$$

By combining these network outputs in parallel, an image of the contents of the vessel is obtained. If the algorithm proceeds to total the seawater, gravel and air pixels in this image, then an estimate of the volume fractions within the cross-section of the vessel can be obtained.

The adjustable parameters, or weights, of the network control the functional mapping and are obtained through network training. Training a neural network is achieved through the presentation of known input–output pairs to the network and adjustment to the network weights until the error is satisfactorily small. To train a neural network requires a large quantity of representative data. Previous research into neural network reconstruction has made use of simulated capacitance and conductance readings [4–9]. Adler and Guardo adopted an interesting approach whereby simulated readings were used to train the network. This network was then tested on a physical impedance tomography system [9]. Good results were achieved provided that the simulated noise was comparable to the measurement noise of the physical system. Since this paper examines the use of neural networks from a practical standpoint, it was decided to train the networks using readings obtained from the physical system. This also eliminated the problem of having to model the noise and drift of the measurement hardware accurately. However, the data would have to be generated in a systematic way in order for it to be feasible to produce the large training database required. It was therefore decided to automate the capturing of data, and




Figure 1. A diagram illustrating individual bubble positions within the cross-section of the vessel. Each circle represents a bubble diameter. The lines form a 10 by 10 matrix of pixels, which are used to set the desired network output.

Table 1. Relative dielectric constants of the phases in the tomography system.

Phase	Relative dielectric constant (at 25 °C)
Air	1.000 54
Polystyrene	2.56
Sand/gravel	3–5
Water	78.54

to develop a system permitting a systematic variation of the contents of the vessel.

The method used for generation of a training database was similar to that employed by Nooralahiyan and Hoyle [6], Nooralahiyan *et al* [5], Hoyle *et al* [4] and Bailey *et al* [7]. Essentially, the smallest identifiable element of one phase, or bubble, is placed in a specific location in the second phase and the readings are taken. The bubble is then moved to a different position and the readings are taken again. This process is repeated until the superposition of all the individual bubble positions effectively covers the cross-section of the vessel. It is important to remember that capacitance tomography systems typically produce low-resolution images. In addition, the readings from the data acquisition system were subject to measurement noise and drift. Consequently, it was decided that a minimum bubble diameter corresponding to 20% of the inner diameter of the vessel would be used. Figure 1 illustrates the way in which these individual bubbles effectively cover the cross-section of the vessel. It is important to note that the term ‘bubble’ does not refer to an air bubble in the usual sense. Instead, it describes a contiguous mass of either gravel or polystyrene.

Air bubbles were represented by polystyrene cylinders. Polystyrene has a dielectric constant that is similar to that of air, as can be seen in table 1, so the capacitance readings were not greatly affected. In addition, it is easier to work with polystyrene. Gravel bubbles were constructed using gravel chips contained in nylon stocking tubes.

The bubble placement system was developed to locate an air or gravel bubble at a specific position within the pipe and seawater volume. It consists of a ‘bed-of-nails’ and clear-acrylic-lid arrangement, as can be seen in figure 2.




Figure 2. A diagram of the bubble placement system showing how the lid and ‘bed-of-nails’ arrangement locates an air bubble at a specific position in the seawater.

Drawn on the clear acrylic lid is a 10 by 10 matrix of pixels used to specify the desired network output. The broken lines in figure 1 outline these pixels. The user can identify which pixels are air, gravel and seawater by looking through the lid from above. A program running on the reconstruction computer provides a graphical interface allowing the user to set the desired phase of each pixel. This program also captures the readings from the data acquisition system. It then saves these readings and an image of the desired network output to a database, which would be used for network training at a later stage.

The training database was generated by placing an individual air bubble in seawater at one of the positions indicated in figure 1 and taking a set of readings. This was also performed for individual gravel bubbles to give 81 sets of readings. Each set of readings was performed three times to improve the noise performance of the neural network. The reasoning is that, for each of these three cases, the desired network outputs would be identical but the voltage readings would differ slightly due to measurement noise. The validation database was generated using combinations of air bubbles, combinations of gravel bubbles and combinations of both air and gravel bubbles. This would determine whether the network was able to generalize from single bubbles of one phase in the training database to multiple bubbles of both phases in the validation database, which clearly represents a more realistic situation.

4. Neural network image reconstruction

This paper examines the use of a multi-layer perceptron and a radial basis function neural network to perform image reconstruction. A multi-layer perceptron is a densely interconnected, layered network of neurons [14]. In particular, this paper will examine the use of a single-layer feed-forward neural network. Each neuron consists of a combination function, which forms the weighted sum of the neuron inputs, and an activation function, which introduces a non-linearity into the network and confines the output to lie within a specific range. The weights of the neuron connections control the functional mapping and are obtained through training. Gradient descent is an iterative process used to modify the weights of the network in order to minimize the error in the mapping. This paper also uses resilient back-propagation to train the networks. Resilient back-propagation is a local, adaptive learning algorithm that uses local gradient

information to modify the weights of the network directly [15]. In both algorithms, the required gradient information is obtained using back-propagation.

Radial basis function neural networks model the non-linear functional mapping between the input and output variables by placing basis functions at specific points in the input space [14]. These basis functions typically take the form of Gaussian kernels that are a function of the Euclidean distance between the input vector and the training data point. The locations of these kernels depend on the specific type of radial basis function neural network being employed. This paper will examine the use of a probabilistic neural network and a kernel adatron.

It is expected that the multi-layer perceptrons will outperform the radial basis function neural networks for this application. This is because radial basis function neural networks are more susceptible to the curse of dimensionality, which specifies that the number of training data points required to ‘fill’ the input space grows exponentially with the dimensions of the input space [14]. Without pre-processing, the neural networks for this application have 56 inputs. Consequently, a large quantity of training data are required in order to obtain an accurate functional mapping. However, the training data are obtained through measurement from a physical system and it is therefore impractical to attempt to generate such a large database.

Neural networks are capable of generalizing for previously unseen bubble configurations. Generalization is described as the ability to predict the correct output for new inputs even though they do not form part of the training database [14]. The ability to generalize is required in the physical system so that the network can perform a reconstruction without the training database having to contain every possible combination of gravel and air bubbles. The generalization capabilities of the networks were tested on a validation database, which represents input combinations not included in the training database, and the following performance measures were used.

- (i) Threshold error, which is the percentage of pixels that are not predicted correctly. For example, if the network predicts that a particular pixel is seawater when it is in fact air, then this represents an error. These errors are summed over the entire database and are divided by the total number of pixels in the database. This measure is an indication of the positional correlation between the desired network output and the network prediction and is intended to assess the accuracy of the reconstructed images. Although it is not the research objective to reconstruct an image of the cross-section of the vessel, the desired volume fraction information can be obtained through analysis of this image.
- (ii) Void fraction error, which is the percentage error of the volume fraction predictions for each of the three phases. For example, if the network predicts that there are 42 pixels of seawater when there are actually 44, then this represents an error of two. These errors are summed over the entire database and are divided by the total number of pixels in the database. This measure is independent of the pixel position and operates simply in terms of volume. It gives an indication of the ability of the network to predict volume fractions accurately.

The generalization capability of the network is tested regularly during training using these performance measures. If the generalization deteriorates because of the network starting to over-fit the training data, then the network training is automatically stopped. This prevents the network from fitting the training data exactly, which would result in a poor reconstruction for unseen bubble configurations.

Before network training commences, the data must be pre-processed. The software standardizes the training database and provides the following pre-processing capabilities.

- (i) Principal component analysis transforms the data so that the inputs are uncorrelated and selects the most important input variables to explain the variance in the target data [14], thereby reducing the input dimensionality. This is likely to improve the generalization performance since fewer weights require training and, therefore, the likelihood of over-fitting is reduced.
- (ii) ‘Removal of dc bias’ ensures that all the inputs for a particular data point have zero mean. This corresponds to removing the dc bias from the output voltages of the sensing circuitry and is an attempt to counteract the measurement hardware drift. The reasoning is that a neural network trained on input voltages from which the dc level has been removed should be insensitive to drift of the sensing circuitry, provided that the circuit does not saturate.

The neural networks are required to perform a classification distinguishing among seawater, gravel and air phases for each pixel. This is a complex task since the dielectric constant of gravel is similar to that of polystyrene, as can be seen in table 1. The results achieved using the multi-layer perceptron will be examined first.

4.1. Multi-layer perceptron image reconstruction

Performing a three-class classification requires a modification to the basic network structure. Nooralahiyan and Hoyle [6] propose a double-sigmoid activation function, which is an extension of the standard sigmoid to provide three distinct output regions. This is achieved through the addition of a low-gradient central region, which permits the network to be trained more readily for target values in this central region. The following activation function structure was decided on for this particular application:

$$f(x) = \frac{1}{\pi}[\tan^{-1}(x - 3) + \tan^{-1}(x + 3)]. \quad (3)$$

Figure 3 is a plot of the double-sigmoid activation function employed. The two threshold values of $\frac{1}{3}$ and $-\frac{1}{3}$ are included to illustrate the three output regions representing air, gravel and seawater, respectively. The target output value for an air pixel is ‘1’, that for a gravel pixel is ‘0’ and that for a seawater pixel is ‘-1’.

A single-layer feed-forward neural network was trained using gradient descent to perform image reconstruction. The network consisted of 56 input neurons corresponding to the 28 capacitance and 28 conductance readings, fully connected to 88 output neurons such that each neuron corresponds to a pixel within the cross-section of the vessel. Figure 4 illustrates




Figure 3. A plot of the double-sigmoid activation function with threshold regions representing air, gravel and seawater, respectively. The independent axis is the weighted neuron input and the dependent axis is the neuron output.




Figure 4. Specific test cases taken from the validation database such that the desired output is on the right-hand side and the prediction of the network is on the left. The white pixels represent air, the dark grey pixels gravel and the light grey pixels seawater. The region of uncertainty between the seawater and air phases has been classified as gravel.

specific test cases taken from the validation database. The desired output is on the right-hand side and the network prediction is on the left-hand side. The specific results achieved will be examined in the next section when a comparison with the 1-of-C output encoding is made.

Although it is simple to implement, we found this technique to be flawed. In general, a multi-class classifier should not be implemented using distinct output levels for each class as is done for the double-sigmoid activation function, since the network performs averaging of outputs. For example, if an input lies between the points with an output of ‘1’ and the points with an output of ‘-1’ in the input space, then the output of the network would be ‘0’. However, this is incorrect since the points with an output of ‘0’ could lie in a completely different region of the input space. In terms of the current application, any pixels for which the output classification is uncertain will be classified as gravel. In particular, air bubbles will have a gravel outline since this represents the region of uncertainty between the seawater and air phases. This situation could be rectified somewhat by interchanging the target values for seawater and gravel. Nooralahiyan and Hoyle [6] also propose various other modifications to improve the reconstruction performance. However, it was decided that a more accurate image reconstruction would be achieved using the standard multi-class classification scheme of 1-of-C encoding.

Table 2. Output encoding for a three-phase tomography system, where y_1 , y_2 and y_3 are the dummy output variables for a specific pixel.

y_1	y_2	y_3	Output phase
1	0	0	Seawater
0	1	0	Gravel
0	0	1	Air

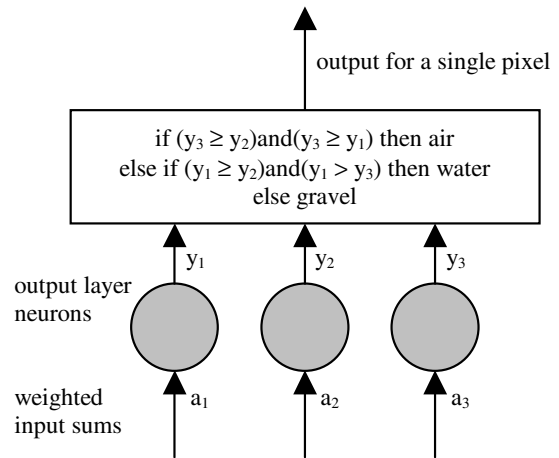


Figure 5. A diagram illustrating 1-of-C encoding on output layer neurons.

A 1-of-C output encoding is performed by introducing additional dummy output variables for each pixel. Each dummy variable is given a target value of zero except for that particular variable corresponding to the desired classification, which has a target value of unity. Table 2 illustrates the encoding scheme for this particular application.

In addition, each output neuron employs a softmax activation function [14]. This ensures that the sum of the three dummy output neurons always equals unity and so these outputs can be interpreted as valid posterior probabilities [14]. The pixel classification is the dummy output neuron with the largest output and, consequently, the greatest posterior probability. The outputs are calculated using the following softmax activation function formula:

$$y_i = \frac{\exp(a_i)}{\sum_{j=1}^3 \exp(a_j)} \quad (4)$$

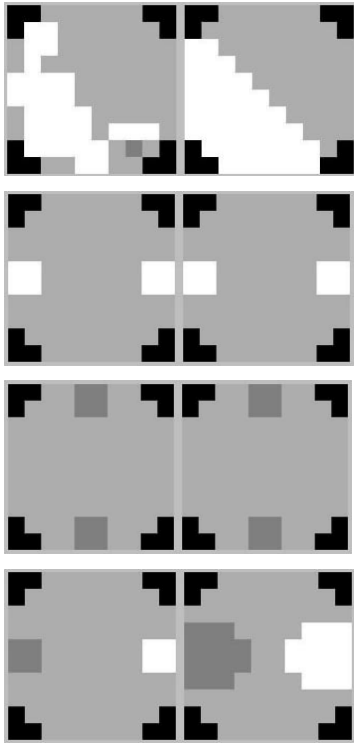
where a is the weighted input for each dummy output neuron. Figure 5 illustrates the phase prediction for a particular pixel within the cross-section of the vessel.

A single-layer feed-forward neural network with a 1-of-C output encoding was trained using resilient back-propagation. The network training results are detailed in table 3, which includes a comparison with the results achieved for the double-sigmoid activation function network. As is evident from table 3, the 1-of-C encoding outperforms the double-sigmoid method in terms of the volume fraction predictions for the air and gravel phases.

Several variations of the data pre-processing were conducted in an attempt to improve the generalization performance of the network. These included removing the dc bias from the input voltages and reducing the input

Table 3. A comparison between the training results of double-sigmoid and 1-of- C output encoded networks.

Performance measure	Double-sigmoid output	1-of- C output encoding
Water void fraction error (%)	11.2163	12.0686
Gravel void fraction error (%)	8.6006	5.6187
Air void fraction error (%)	11.0880	9.1183
Validation threshold error (%)	21.7593	19.9453

**Figure 6.** Specific test cases taken from the validation database for a network with 1-of- C output encoding. The desired network output is on the right-hand side and the prediction of the network is on the left-hand side. The white pixels represent air, the dark grey pixels gravel and the light grey pixels seawater.

dimensionality using principal component analysis, whereby only those inputs that contribute more than 1% to the output variations were retained. This resulted in only ten of the original 56 inputs being retained. The network training results are detailed in table 4. As is evident from table 4, the removal of the dc bias improved the overall accuracy of the volume fraction predictions.

Figure 6 illustrates specific test cases taken from the validation database for a network from which the dc bias has been removed. 1-of- C output encoding was employed and the network was trained using resilient back-propagation. The desired network output is on the right-hand side and the prediction of the network is on the left-hand side.

These test cases indicate the ability of the network to discriminate between air and gravel bubbles. They also indicate the ability of the network to generalize from single bubbles of either air or gravel in the training database to combinations of both air and gravel bubbles in the validation database. This is a promising result since the network was

not trained with any bubble configurations in which all three phases were included simultaneously.

According to Bishop [14], multi-layer perceptrons with sigmoidal activation functions and two layers of adaptive weights can approximate any decision boundary to arbitrary accuracy. A double-layer feed-forward neural network was implemented to determine whether any improvement in the image reconstruction could be achieved. The optimal number of neurons in the hidden layer was found using a model search algorithm and the resulting network was trained using resilient back-propagation. However, the addition of a hidden layer did not improve the accuracy of the reconstruction, confirming the theory of Bailey *et al* [7], who concluded that the image reconstruction for a twelve-electrode capacitance tomography system demonstrates linear separability. A single-layer feed-forward neural network is able to construct a separating hyperplane. Consequently, the addition of a hidden layer of neurons will not improve the performance of the network. It is also likely that insufficient training data will be available to model the non-linearity of the functional mapping accurately.

4.2. Radial basis function neural network image reconstruction

Probabilistic neural networks and kernel adatoms were employed to perform image reconstruction. A principal component analysis was performed for both networks to reduce the input dimensionality from 56 to ten and the network training results are detailed in table 5. Table 5 also includes a comparison with the single-layer feed-forward neural network discussed earlier.

As is evident from table 5, the performances of the radial basis function neural networks were inferior to that of the single-layer feed-forward neural network. This was expected since the limited training database is insufficient to ‘fill’ the input space. Consequently, the radial basis function neural networks are unable to construct an accurate mapping due to the curse of dimensionality. The following section examines a simple modification to the network structure that improves the accuracy of the volume fraction predictions.

5. The neural network volume fraction predictor

In terms of the intended application, accurate volume fraction predictions are required. Previous neural network reconstruction techniques have concentrated on producing an image of the cross-section of the vessel. The volume fractions of the phases are then obtained through analysis of these images. This section examines an alternative technique whereby the network is required to predict the volume fractions directly rather than by performing image reconstruction.

The neural network is required to model the functional mapping between the 56 readings obtained from the data acquisition system and the continuous outputs representing the volume fractions of the seawater, gravel and air phases. Modifications to the network structure are therefore required. Firstly, the number of output neurons is reduced to two, giving the number of pixels of gravel and the number of pixels of air, respectively. The number of pixels of seawater is found by virtue of the fact that 88 pixels cover the cross-section of

Table 4. A comparison between the networks with modified pre-processing and the original 1-of-C output encoded network.

Performance measure	Original network	PCA pre-processing	Removal of DC bias
Water void fraction error (%)	12.0686	13.7584	11.1890
Gravel void fraction error (%)	5.6187	5.9870	6.1553
Air void fraction error (%)	9.1183	10.9701	9.3224
Validation threshold error (%)	19.9453	19.6170	17.8620

Table 5. A comparison between the radial basis function neural networks and the single-layer feed-forward neural network discussed earlier.

Performance measure	Single-layer feed-forward	Probabilistic neural network	Kernel adatron
Water void fraction error (%)	12.0686	19.8106	16.2310
Gravel void fraction error (%)	5.6187	6.5909	5.6629
Air void fraction error (%)	9.1183	13.5985	13.1439
Validation threshold error (%)	19.9453	21.6288	21.3068

Table 6. A comparison between the volume fraction predictor and the best image reconstruction neural network.

Performance measure	Image reconstruction neural network	Volume fraction predictor
Water void fraction error (%)	11.1890	6.4843
Gravel void fraction error (%)	6.1553	4.9868
Air void fraction error (%)	9.3224	8.9478

the vessel. Secondly, linear activation functions were used for these two output neurons. Consequently, the neural network performs a regression function, namely

$$[\text{gravel}_{\text{FRACTION}} \text{air}_{\text{FRACTION}}] = f_{\text{REGRESSION}}(\vec{C}, \vec{R}) \quad (5)$$

where

$$\left. \begin{array}{l} \text{gravel}_{\text{FRACTION}} \\ \text{air}_{\text{FRACTION}} \end{array} \right\} \in [0, 88]. \quad (6)$$

A double-layer feed-forward neural network was trained using resilient back-propagation to perform this task. The network training results are detailed in table 6, which includes a comparison with the void fraction results achieved for the best image reconstruction network.

As is evident from table 6, the volume fraction predictor outperforms the image reconstruction neural network, indicating that more accurate volume fraction predictions can be obtained by ignoring the task of image reconstruction and predicting the volume fractions directly. This is likely to be a result of there being fewer network weights since there are only two outputs instead of the original 264. Given the limited training data available, a more accurate model can be obtained if fewer network weights require training.

6. Conclusion

This paper illustrates that it is feasible to use a neural network as an alternative image reconstruction tool for impedance tomography systems. In particular, the network was able to discriminate among the seawater, gravel and air phases and generalize to previously unseen bubble configurations. It was also shown that the accuracy of the volume fraction predictions could be improved by training a neural network to predict the volume fractions of the various components directly.

References

- [1] Tapson J 1999 Neural networks and stochastic search methods applied to industrial capacitance tomography *Control Eng. Practice* **7** 117–21
- [2] Pinheiro P A T, Loh W W and Dicken F J 1998 Optimal sized electrodes for electrical resistance tomography *Electron. Lett.* **34** 69
- [3] Dicken F J, Williams R A and Beck M S 1993 Determination of composition and motion of multicomponent mixtures in process vessels using electrical impedance tomography—I. Principles and process engineering applications *Chem. Eng. Sci.* **48** 1883–97
- [4] Hoyle B S, Bailey N J and Nooralahiyan A Y 1995 Performance of neural network in capacitance-based tomographic process measurement systems *Meas. Control* **28** 109–12
- [5] Nooralahiyan A Y, Hoyle B S and Bailey N J 1994 Neural network for pattern association in electrical capacitance tomography *Proc. IEE* **141** 517–21
- [6] Nooralahiyan A Y and Hoyle B S 1997 Three-component tomographic flow imaging using artificial neural network reconstruction *Chem. Eng. Sci.* **52** 2139–48
- [7] Bailey N J, Hoyle B S, Spink D M and Nooralahiyan A 1998 Neural networks in process tomography <http://www.ee.leeds.ac.uk/homes/NJB/Research/bergen/bergen.html>
- [8] Ratajewicz-Mikolajczak E, Shirkoohi G H and Sikora J 1998 Two ANN reconstruction methods for electrical impedance tomography *IEEE Trans. Magn.* **34** 2964–7
- [9] Adler A and Guardo R 1994 A neural network image reconstruction technique for electrical impedance tomography *IEEE Trans. Med. Imag.* **13** 594–60
- [10] Marioli D, Sardini E and Taroni A 1993 High-accuracy measurement techniques for capacitance transducers *Meas. Sci. Technol.* **4** 337–43
- [11] Marioli D, Sardini E and Taroni A 1991 Measurement of small capacitance variations *IEEE Trans. Instrum. Meas.* **40** 426–8
- [12] Smit Q, Tapson J and Mortimer B J P 2000 Material phase detection system using capacitance tomography *Proc. IEEE-IMTC Conf. (Baltimore, WA)*
- [13] Stott A L, Green R G and Seraji K 1985 Comparison of the use of internal and external electrodes for the measurement of the capacitance and conductance of fluids in pipes *J. Phys. E: Sci. Instrum.* **18** 587–92
- [14] Bishop C M 1995 *Neural Networks for Pattern Recognition* (New York: Oxford University Press)
- [15] Riedmiller M and Braun H 1993 A direct adaptive method for faster backpropagation learning: the RPROP algorithm *Proc. IEEE Int. Conf. on Neural Networks (ICNN) (San Francisco)* ed H Ruspini, pp 586–91