

APPLICATIONS OF NEURAL NETWORKS TO ULTRASOUND TOMOGRAPHY

B. C. Conrath, C. M. W. Daft, and W. D. O'Brien, Jr.

Department of Electrical and Computer Engineering
University of Illinois
Urbana, IL 61801

ABSTRACT

Ultrasound tomography holds promise in the area of medical diagnosis but is limited by the inadequacies of current reconstruction algorithms. This paper presents an alternative method using neural networks. The theory presented begins with X-ray tomography and then extends to ultrasound. After providing theoretical relevance, several experiments are discussed, illustrating the effectiveness of neural networks. The model was a circular cylinder with acoustic properties of tissue insonated by a line source at 2 MHz. The transducers were arranged in a ring surrounding the cylinder with one being the transmitter. The experiments involved varying both the acoustic speed and the radius of the cylinder. In both cases, the neural network was able to generalize to parameters other than the ones used during the training.

I. INTRODUCTION

Neural networks are naturally suited to perform the inversion of tomography data. Carroll and Dickinson have shown how a feedforward neural network can implement the filtered backprojection algorithm [1]. In fact, neural networks can, in principle, "learn" any function [2]. However, important practical issues remain, such as the morphology of an efficient neural network. Conventional inversion methods using the Fourier Diffraction Theorem assume weak scattering, resulting in a loss of accuracy. Frequency domain interpolation is prone to instability problems. The approach adopted in this paper, using a neural network, has been shown experimentally to be an effective method, and promises to offer a better alternative. In order to illustrate how neural networks can be utilized, a natural evolution was traced from the X-ray tomography case to the ultrasound case. Experiments with neural networks follow, showing that this method is capable of reconstructing characteristics of an object with good accuracy, given scattered field data.

II. MOTIVATION

X-ray tomography possesses a computationally efficient inversion algorithm known as filtered backprojection. Using the fan-beam geometry (figure 1), the filtered backprojection algorithm gives the following object function $f(r, \phi)$

$$f(r, \phi) = \int_0^{2\pi} \frac{1}{L^2} Q_{\beta}(\gamma) d\beta \quad (1)$$

where

$$L = L(r, \phi, \beta) = \sqrt{[D+r\sin(\beta-\phi)]^2 + [r\cos(\beta-\phi)]^2} \quad (2)$$

and

$$\gamma = \tan^{-1} \frac{r\cos(\beta-\phi)}{D+r\sin(\beta-\phi)} \quad (3)$$

$Q_{\beta}(\gamma)$ is the filtered backprojection information at an angle of incidence β [3]. $Q_{\beta}(\gamma)$ can be found by using two different methods (figure 2). Using the convolution route, $Q_{\beta}(\gamma)$ becomes

$$Q_{\beta}(\gamma) = P'_{\beta}(\gamma) \otimes g(\gamma) \quad (4)$$

where

$$P'_{\beta}(\gamma) = D P_{\beta}(\gamma) \cos \gamma \quad (5)$$

$P_{\beta}(\gamma)$ is the raw projection data. In order to implement (1) using a linear, parallel distributed processing (PDP) network, the object function equations were discretized. Assuming that the location of interest (x, y) was situated upon the line between the transmitter and a receiver, γ can be matched to $n\alpha$, where α denotes the angle increment between adjacent transducers. Now $f(x, y)$ can be approximated by

$$f(x, y) \cong \frac{2\pi}{M} \sum_{i=1}^M \frac{1}{L^2(x, y, \beta[i])} Q_{\beta[i]}(n\alpha) \quad (6)$$

where

$$Q_{\beta[i]}(n\alpha) = \alpha \sum_{k=-(N-1)}^{N-1} g(n\alpha - k\alpha) P'_{\beta[i]}(k\alpha) \quad (7)$$

for $n = 0, 1, 2, \dots, N-1$

A linear network to implement these equations is shown in figure 3. This network implements (6) in the following way:

- 1) Select an appropriate point of interest within the object (x,y).
- 2) Irradiate the object at M equally spaced angles $\beta[i]$, $i=1$ to M.
- 3) Gather projection information $P_{\beta[i]}(\gamma)$.
- 4) Sample the projection at discrete angle intervals α to obtain $P_{\beta}(n\alpha)$.
- 5) Modify the projection $P'_{\beta}(n\alpha)$.
- 6) Convolve the modified projection with the filter $g(n\alpha)$ where

$$g(n\alpha) = \begin{cases} \frac{1}{8\alpha^2} & n=0 \\ 0 & n \text{ even} \\ \left(\frac{\alpha}{\pi\alpha\sin(n\alpha)}\right)^2 & n \text{ odd} \end{cases} \quad (8)$$

The sampled modified projection serves as the input to the network. The connection weights contain the filter information shifted to perform the convolution. The network multiplies the input by the value of the connection weight, and the weighted signal is then passed on to the summation function. This process is repeated in the next layer. Therefore, the network (figure 3) realizes the object function at a point (x,y) for any object.

This theory can be extended to the ultrasound case, considering first only the diffraction effects. In the limit of high frequencies, and in the absence of refraction, the acoustic and X-ray problems become analogous, with acoustic attenuation replacing density. Starting with this linear network at a high frequency, a smooth transition to the ultrasound diffraction case can be made by gradually lowering the frequency and introducing a nonlinearity after the network's summation functions. This nonlinearity, a characteristic property of neural networks, broadens the range of functions the network can "learn." In this case, the nonlinearity is a hyperbolic tangent. After slightly reducing the frequency, the network can be trained again. This gentle process can be repeated until a network has been trained that compensates for diffraction and avoids local minima in weight space. (See Section III for details.) The linear network can solve the acoustic problem for simple geometries when the object function has a linear effect on the scattered field,

because the diffraction effects can be calculated and included in the weights. A neural network is necessary if the geometry is complex or unknown, as is always the case with tissue.

These results can be extended to acoustic speed reconstruction, by using phase information. Again, starting from the X-ray case, assume that the phase information can be collected. The real and imaginary parts of the object function can be calculated separately using two distinct linear networks. The ultrasound case requires that the networks for the real and imaginary parts be interconnected. The weights of the two distinct networks in the X-ray case can be set a priori. The weights on the connections extending between the real and imaginary networks can be set to very small, random values. The neural network can then fine-tune the weights to model the ultrasound case accurately. As a preliminary exercise showing that the neural network is capable of inverting scattered field data, we have trained a network for a simple geometry.

III. EXPERIMENTS USING NEURAL NETWORKS

The previous discussion introduced the idea of a linear network representing the filtered backprojection equations for fan-beam X-ray tomography. The ultrasound case is made more difficult due to diffraction and refraction effects. The neural network with a morphology inspired by the linear network can overcome these difficulties. Instead of entering the value of the connection weights, the neural network can learn the correct weights when repeatedly presented with example input-output pairs.

A backpropagation neural network was used in the following experiments (figure 4). The network was constructed with three fully connected layers and a bias. The connections between the processing elements hold weights, values multiplied by the signal passing through the connection. When the weighted signals reach the processing element, they are added together internally before passing the resultant signal through a nonlinearity, typically a sigmoid function, but hyperbolic tangents and sinusoids are also possible (figure 5). The process can be represented by

$$x_j^{[s]} = f\left(\sum_{i=0}^n [w_{ji}^{[s]} x_i^{[s-1]}]\right) = f(I_j^{[s]}) \quad (9)$$

and

$$f(I_j^{[s]}) = \frac{1}{1 + \exp(-I_j^{[s]})} \quad \text{for sigmoid.} \quad (10)$$

The layers are indexed by s, s-1, etc.

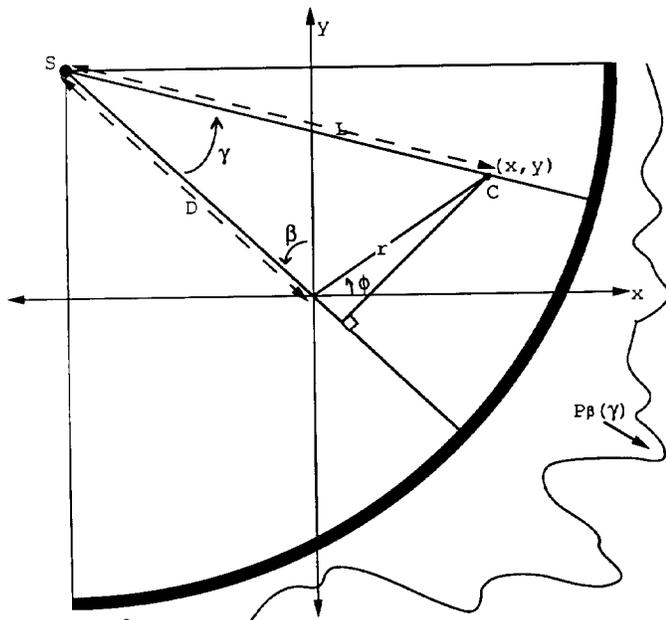


FIGURE 1. FAN-BEAM GEOMETRY

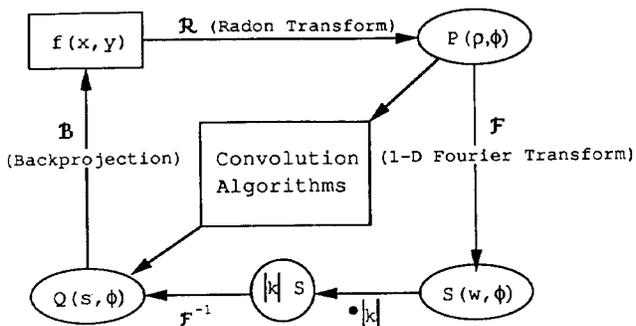


FIGURE 2. TWO METHODS FOR FILTERED BACKPROJECTION

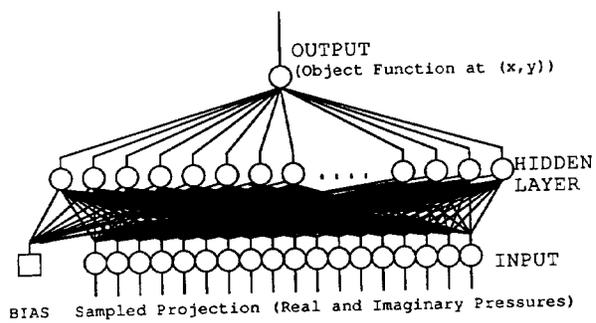


FIGURE 4. BACKPROPAGATION NEURAL NETWORK

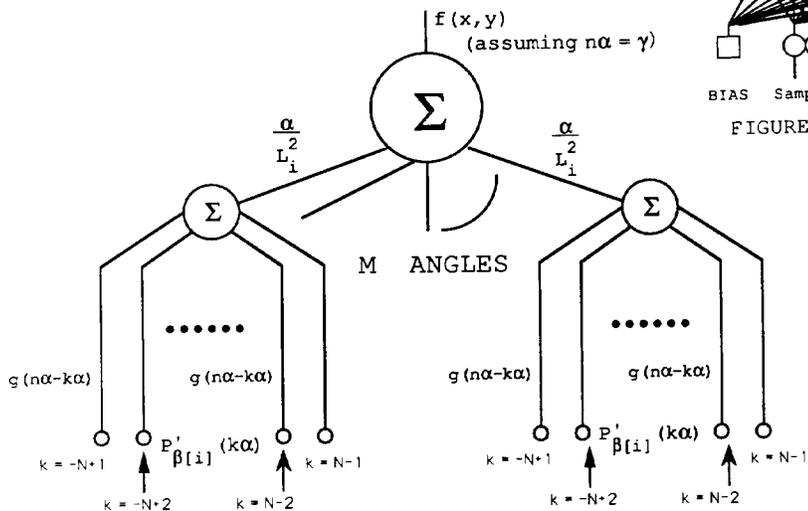


FIGURE 3. LINEAR NETWORK

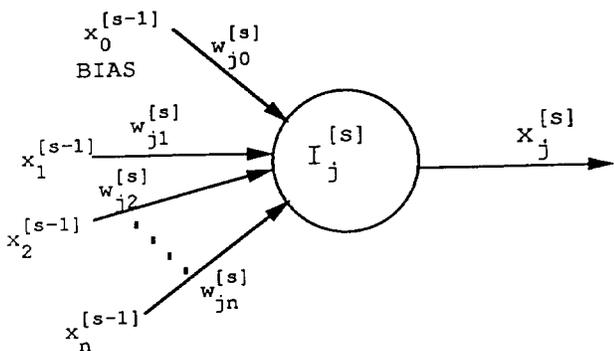


FIGURE 5. PROCESSING ELEMENT

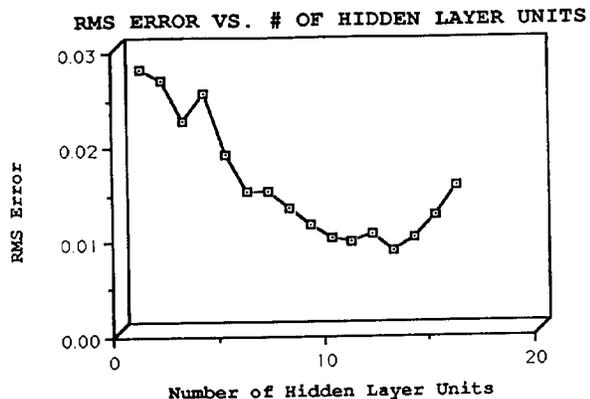


FIGURE 6. RMS ERROR FOR RADIUS EXPERIMENT

The weights are usually initialized to a random value. Training the network involves altering the weights to a correct value that gives the desired output. Training begins by presenting an input i to each of the elements in the input layer. These values propagate through the network to the output layer in according to (9) and (10). At the output layer, the local error e is calculated by comparing the network output o with the desired output d .

$$e_k = d_k - o_k \quad (11)$$

where k is the index of the output processing elements. The connection weights w are then altered according to

$$\Delta w_{ji}^{[s]}(t) = C_1 e_j^{[s]} x_i^{[s-1]} + C_2 \Delta w_{ji}^{[s]}(t-1) \quad (12)$$

where $\Delta w_{ji}^{[s]}(t-1)$ is the previous weight

change, and C_1 and C_2 represent learning coefficients. The local error and the delta weight are calculated by propagating back through each layer. Then the weights are adjusted by adding the corresponding weight change. This process controls learning.

The experimental model of tissue was a circular cylinder insonated by a 2 MHz line source. A ring of 16 transducers, with one acting as the transmitter, surrounded the cylinder. The radius of the transducer ring was 75 mm. The medium surrounding the cylinder had the acoustic properties of water. The first experiment used a non-attenuating cylinder with a radius of 0.5625 mm. The network was trained with speeds varying from 1526 to 1625 ms^{-1} . Although a neural network was not necessary, a network with just one hidden layer unit gave the best accuracy, confirming the simplicity of this problem.

The second experiment involved varying the radius from 0.454 mm to 0.754 mm, while keeping the speed constant at 1576 ms^{-1} . The neural network was necessary for this kind of problem. The networks were trained with different numbers of elements in the hidden layer, using a training set of 80 input-output pairs presented 20,000 times. The test set consisted of 80 random radii within the training range. The input was the real and imaginary pressures from the transducers, and the output was the radius. The neural network with 13 hidden layer units learned with the greatest accuracy with an RMS error of 0.00829 (figure 6). Networks with more than 16 hidden units failed to learn after the 20,000 presentations, suggesting that larger networks require more presentations and have a greater likelihood of falling into local minima in weight space. Neural networks can learn to calculate the radius of an object from its scattered field data, and could possibly achieve greater accuracy by

adjusting parameters, such as the learning rate, initial weight values, scaling of the input data, and size of the hidden layer.

IV. CONCLUSION

Neural networks can "learn" characteristics about a scanned object from its projection data. More research is necessary to determine an efficient morphology and learning rate for specific problems. Neural networks trained from completely randomized weights require many training cycles. Starting with an acoustic analog to the X-ray approximation can possibly decrease the training time. A neural network to analyze an entire image is very large. For instance, a 100 x 100 pixel image requires 10,000 interconnected networks that input both real and imaginary pressures. Assume that 100 samples were taken from every one of 100 projections, so that the input layer of each network contains 10,000 elements. The entire input layer contains on the order of 10^8 elements. The hidden layer contains on the order of 10^6 elements, and the output layer contains 10,000 elements. Although this network is very large, neural network hardware, using optical or VLSI implementations, will probably be able to perform a problem of this magnitude. This method does promise to provide greater accuracy, because the neural network does not assume weak scattering. The neural network can also be fine-tuned, according to instrument-dependent variables.

Further experiments will include training the network conceived in Section II and examining the effects of various morphologies on accuracy. Starting the training using weights from the X-ray case instead of random weights will also be examined.

V. ACKNOWLEDGEMENTS

We are grateful to Dr. Y. Bresler for helpful discussions. This research was supported by the Whitaker Foundation, the National Institutes of Health (CA 36029), and the University of Illinois Campus Research Board.

REFERENCES

- [1] S. M. Carroll and B. W. Dickinson, "Construction of Neural Nets Using the Radon Transform," Proc. IJCNN Vol. 1, 1989, pp. 607-611.
- [2] G. Cybenko, "Approximations by Superpositions of a Sigmoidal Function," to appear in Mathematics of Control, Signals, and Systems, 1989.
- [3] A. C. Kak and M. Slaney, Principles of Computerized Tomographic Imaging (New York: IEEE Press, 1988), pp. 49-86.