

Super-Resolution Target Identification from Remotely Sensed Images Using a Hopfield Neural Network

Andrew J. Tatem, Hugh G. Lewis, Peter M. Atkinson, and Mark S. Nixon

Abstract—Fuzzy classification techniques have been developed recently to estimate the class composition of image pixels, but their output provides no indication of how these classes are distributed spatially within the instantaneous field of view represented by the pixel. As such, while the accuracy of land cover target identification has been improved using fuzzy classification, it remains for robust techniques that provide better spatial representation of land cover to be developed. Such techniques could provide more accurate land cover metrics for determining social or environmental policy, for example. The use of a Hopfield neural network to map the spatial distribution of classes more reliably using prior information of pixel composition determined from fuzzy classification was investigated. An approach was adopted that used the output from a fuzzy classification to constrain a Hopfield neural network formulated as an energy minimization tool. The network converges to a minimum of an energy function, defined as a goal and several constraints. Extracting the spatial distribution of target class components within each pixel was, therefore, formulated as a constraint satisfaction problem with an optimal solution determined by the minimum of the energy function. This energy minimum represents a “best guess” map of the spatial distribution of class components in each pixel. The technique was applied to both synthetic and simulated Landsat TM imagery, and the resultant maps provided an accurate and improved representation of the land covers studied, with root mean square errors (RMSEs) for Landsat imagery of the order of 0.09 pixels in the new fine resolution image recorded. As such, we show how, by using a Hopfield neural network, more accurate measures of land cover targets can be obtained compared with those determined using the proportion images alone. The Hopfield neural network used in this way represents a simple, robust, and efficient technique, and results suggest that it is a useful tool for identifying land cover targets from remotely sensed imagery at the subpixel scale.

Index Terms—Fuzzy image classification, Hopfield networks, image resolution, land cover, optimization methods, super-resolution object detection.

I. INTRODUCTION

INFORMATION on land cover features is required for management and understanding of the environment. Accurate identification and extraction of target land cover features is a vital procedure for many areas of work, e.g., military intelligence, agricultural planning, and water resource management. Remote sensing has the potential to provide this information. Imagery derived from aircraft and satellite-mounted sensors

have attributes that make remote sensing suitable for target identification, and the short orbit times of many sensors (of the order of 3 h to 16 days) means a temporal sequence of images can be acquired, aiding the monitoring of specific features. Also, the large spatial coverage that can be obtained with imagery from satellite sensors (e.g., of the order of 1×10^6 km²) provides an advantage over costly and time-consuming ground survey. Many remote sensors measure ground reflectance at a fine spectral resolution and for most target identification applications, this provides sufficient information to identify accurately features of interest. Finally, remote sensing has the potential to provide land cover target information at a variety of scales (e.g., from < 1 m to 1 km). These aspects make target identification from remotely sensed imagery attractive. However, there exist several practical limitations.

Perhaps the biggest drawback of target identification from remotely sensed images relates to that of scale. Spatial scale is a key factor in the interpretation of remotely sensed land cover data [1], and the information obtainable from such imagery can vary greatly depending on the spatial variation in the observed land cover and the specific terrain characteristics under consideration. There also exist practical limits to the level of detail that can be identified by each remote sensor and these limits are defined by the resolutions of the remote sensing system. One of the commonest measures of image characteristic used is spatial resolution, which determines the level of spatial detail depicted in an image. This measure is a function of the instantaneous field-of-view (IFOV) of a sensor, defined as the cone angle within which incident energy is focused on the detectors [2]. In turn, the IFOV leads to a ground resolution element (GRE) on the surface of the Earth (this GRE should not be confused with the pixel, which is the output product to which a radiance value is assigned).

The pixel represents the smallest element of a digital image and has, therefore, traditionally represented a limit to the spatial detail obtainable in target feature extractions from remotely sensed imagery. Within remotely sensed images, a significant proportion of pixels is often of mixed land cover class composition, and their presence can adversely affect the performance of image analysis operations [3]. The solution to the mixed pixel problem typically centres on fuzzy classification. Subpixel class composition is estimated through the use of techniques such as spectral mixture modeling [4], multilayer perceptrons [5], nearest neighbor classifiers [6], and support vector machines [7]. These approaches allow proportions of each pixel to be partitioned between classes, and the output of these techniques generally takes the form of a set of proportion images, each displaying the proportion of a certain class within each pixel. In

Manuscript received December 19, 1999; revised July 10, 2000. This work was supported by EPSRC studentship 98 321 498.

A. J. Tatem, H. G. Lewis, and M. S. Nixon are with the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K. (e-mail: A.J.Tatem@soton.ac.uk).

P. M. Atkinson is with the Department of Geography, University of Southampton, Southampton, U.K.

Publisher Item Identifier S 0196-2892(01)02140-4.

most cases, this results in a more appropriate and informative representation of targets than that produced using a hard, one class per-pixel classification. However, while the class composition of every pixel is estimated, the spatial distribution of these class components within the pixel remains unknown.

The work in this paper, along with other work in the literature, demonstrates that it is possible to identify land cover targets at the subpixel scale (super-resolution). Fisher [3] argues that within remotely sensed images, the pixel can only be subdivided by creating repetitive information and Fig. 1 demonstrates this argument, with the spatial resolution of a 2×2 pixel image containing four classes (A, B, C and D) being increased to a 4×4 pixel image. The finer resolution means that all new pixels fit within the old pixels and take the same value. However, the technique described in this paper suggests that the use of prior information allows greater precision. Such a land cover target identification technique potentially leads to several useful applications.

- 1) **To identify land cover targets at a fine spatial resolution from any remote sensing system.** While sensors on satellites such as IKONOS (up to 1 m spatial resolution) can provide sufficient spatial detail for accurate land cover target extraction, the cost and availability of such data may prohibit its use in many areas of work. By applying the developed technique to cheaper, more readily available data, for example SPOT HRV (up to 10 m spatial resolution) data, similar levels of accuracy for land cover target extraction might be achieved.
- 2) **To apply such a technique to obtain more accurate land cover metrics from remotely sensed imagery.** By increasing the spatial resolution of land cover target maps derived from medium-resolution sensors such as the Landsat thematic mapper (TM) or the SPOT high resolution visible sensor, the potential exists to, for example, more accurately locate field boundaries or define areas of semi-natural vegetation. Such information would be of use in determining environmental or social policy, for example.
- 3) **For fine detail urban target identification.** With the advent of satellites such as IKONOS and Orbview, and the more common use of airborne remote sensing, imagery of spatial resolution less than 5 m is becoming widely available. Application of a technique to produce super-resolution maps of land cover targets from these source data would allow urban land cover target extraction and mapping of an unprecedented fine detail from remotely sensed imagery.
- 4) **To simulate fine spatial resolution imagery from imagery of a coarser spatial resolution.** Such an approach could aid decision making on future choices of imagery.

A. Previous Work

Only recently has research been undertaken on the subject of identifying land cover targets from remotely sensed images at the subpixel scale. Schneider [8] introduced a knowledge-based analysis technique for the automatic localization of field boundaries with subpixel accuracy. The technique relies on knowledge



Fig. 1. Example of Fisher's argument.

of straight boundary features within Landsat TM scenes, and serves as a preprocessing step prior to automatic pixel-by-pixel land cover classification. With knowledge of pure pixel values either side of a boundary, a model is defined for each 3 by 3 block of pixels, the model then uses parameters such as pure pixel values, boundary angle, and distance of boundary from the centre pixel. Using least squares adjustment, the most appropriate parameters are chosen to locate a subpixel boundary that divides mixed pixels into their respective pure components. Improvements on this technique were described in [9] which used a neural network to speed up processing and [10], [11] and [12] put forward algorithmic improvements, along with the addition of a vector segmentation step. The technique represents a successful, automated and simplistic pre-processing step for increasing the spatial resolution of satellite imagery. However, its application is limited to features with straight boundaries at a certain spatial resolution and the models used still have problems resolving image pixels containing more than two classes [13].

Flack *et al.* [14] also concentrated on super-resolution target identification at the borders of agricultural fields, where pixels of mixed class composition occur. Edge detection and segmentation techniques were used to identify field boundaries and the Hough transform [15] was applied to identify the straight, subpixel boundaries. These vector boundaries were superimposed on a subsampled version of the image, and the mixed pixels were reassigned each side of the boundaries. By altering the image subsampling, the degree to which the spatial resolution was increased could be controlled. However, no validation or further work was carried out, and so the success of the technique remains unclear.

Aplin *et al.* [16] also made use of subpixel scale vector boundary information, along with fine resolution satellite sensor imagery to identify land cover targets. By utilizing Ordnance Survey land line vector data and undertaking per-field rather than the traditional per-pixel land cover classification, target identification at a subpixel scale was demonstrated. Assessments suggested that the per-field classification technique was generally more accurate than the per-pixel classification. However, in most cases around the world, accurate vector data sets with which to apply the approach will rarely be available.

The techniques described so far are based on direct processing of the raw imagery. Other work has focused on using a preprocessing step where fuzzy classification of the imagery is undertaken, and an attempt to map the location of class components within the pixels is made.

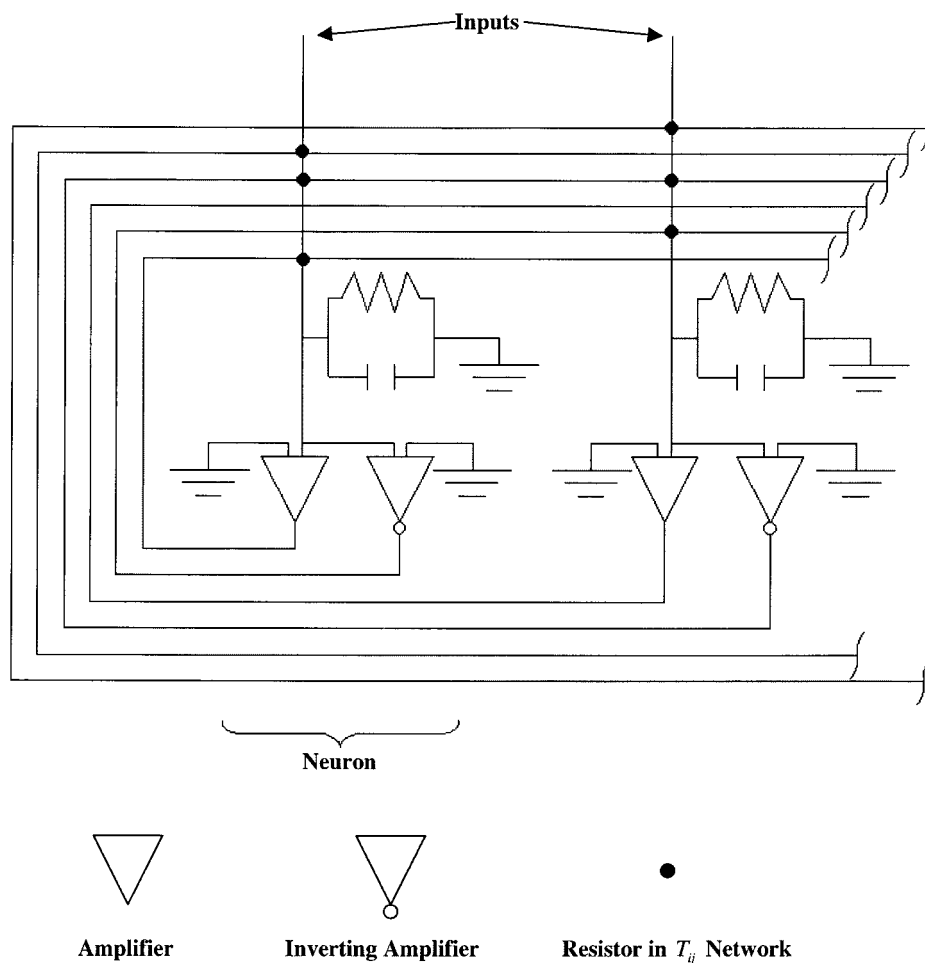


Fig. 2. Hopfield neural network as an analog circuit. The black circles at the intersections represent resistive connections (T_{ij} s) between outputs and inputs. Connections between inverted outputs and inputs represent negative connections.

Atkinson [17] used an assumption of spatial order within and between pixels to map the location within each pixel of the proportions output from a fuzzy classification. Spatial dependence is the concept whereby observations close together tend to be more alike than those further apart [18], and this assumption proved to be valid for recreating the spatial distribution and areal coverage of the land cover. The algorithm produced a certain degree of success for semi-natural land cover targets. However, the simple technique suffered from problems due to the complex land cover mixing in the data used.

Unlike Atkinson [17], Foody [19] made use of additional information in the form of a second higher spatial resolution image. The author used it in a simple regression based approach to sharpen the output of a fuzzy classification of a lower spatial resolution image, producing a subpixel land cover map. The results produced a visually improved representation of the lake being studied, and this was further improved by fitting class membership contours, lessening the blocky nature of the representation. However, the areal extent of the lake was not maintained using the contouring technique and in addition, the author noted that it is difficult to obtain two coincident images of differing spatial resolution.

B. Paper Structure

This paper describes an approach that uses the output from a fuzzy classification technique to constrain a Hopfield neural network formulated as an energy minimization tool. In Section II, an overview of the workings and previous applications of the Hopfield neural network will be given. Section III introduces the modifications made to apply the Hopfield network to the problem of land cover target identification and mapping at the subpixel scale, as well as methods used to understand and improve the processes at work. In Section IV, results of applying the approach to synthetic imagery are used to explain and understand the performance of the network. Section V illustrates the results of applying the technique to simulated remotely sensed imagery, and Section VI provides analysis and explanation of these results. Finally, Section VII provides a summary and conclusion of the findings of the research.

II. HOPFIELD NEURAL NETWORK

The Hopfield neural network is a fully connected recurrent network and can be implemented physically by interconnecting a set of resistors and amplifiers with symmetrical outputs and external bias current sources (Fig. 2). The mathematical model

describing the behavior of such an array of electronic components can be derived from Kirchoff's current law [20], [21]

$$C_i \frac{du_i}{dt} = -\frac{u_i}{R_i} + \sum_{j=1}^N T_{ij} v_j + I_i; \quad i = 1 \dots N \quad (1)$$

where

$$\frac{1}{R_i} = \frac{1}{R_{i0}} + \sum_{j=1}^N \left(\frac{1}{R_{ij}^+} + \frac{1}{R_{ij}^-} \right). \quad (2)$$

R_{ij} is the resistance between the output of amplifier j and input of amplifier, i , N is the number of amplifiers, $C_i > 0$ is the capacitance of amplifier, i , u_i is the internal voltage of amplifier i , and I_i is the external bias on amplifier i . T_{ij} is the conductance from amplifier j to amplifier i , where

$$T_{ij} = \frac{1}{R_{ij}^+} + \frac{1}{R_{ij}^-} \quad (3)$$

and $v_i = g_i(u_i)$ is the output voltage of amplifier i . $g_i(u_i)$ is the nonlinear activation function, defined as

$$g_i(u_i) = \frac{1}{2}(1 + \tanh \lambda u_i) \quad (4)$$

where λ determines the steepness of the function.

Hopfield [22] shows how (1) can be written in a neural context for ease of interpretation, where, in this case the nonlinear amplifiers correspond to neurons

$$\tau_i \frac{du_i}{dt} = -\alpha_i u_i + \sum_{j=1}^N T_{ij} v_j + I_i; \quad i = 1 \dots N \quad (5)$$

where

- $\tau_i = C_i$ time constant for neuron i ;
- u_i total weighted input at neuron i , $\alpha_i = (1/R)_i$;
- N number of neurons in the network;
- $v_i = g_i(u_i)$ neural output which is a function of the input u_i ;
- I_i external bias on neuron i ;
- T_{ij} weight from neuron j to neuron i .

which corresponds to the conductance in (1).

The set of differential equations described so far defines the time evolution of the network. Thus, from a set of initial neuron outputs, the state v of the network varies with time until convergence to a stable state, where neuron output stops varying with time. Weights and biases determine the neural outputs at this stable state.

Hopfield [22] showed that using symmetric weights with no self connection, i.e., $T_{ji} = T_{ij}$ and $T_{ii} = 0$ is sufficient to guarantee convergence to such a stable state. Therefore, independent of its initial status, a Hopfield neural network will always reach an equilibrium state where no output variation occurs and it was also demonstrated that for high values of the gain, λ_i , the activation function $g_i(u_i)$ (4) approaches a step function. The

stable states of the network consequently correspond to the local minima of the following energy function [21]:

$$E = -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N T_{ij} v_i v_j - \sum_{i=1}^N v_i I_i + \sum_{i=1}^N \frac{\alpha_i}{\lambda} \int_0^{v_i} g_i^{-1}(v) dv \quad (6)$$

where E is the energy calculated over the whole network. For neurons where $\alpha \approx 0$ and with high values of λ , the last term becomes small and can be neglected, such that

$$E = -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N T_{ij} v_i v_j - \sum_{i=1}^N v_i I_i. \quad (7)$$

From (5) and (7), the equation describing the dynamics, i.e., the rate of change of neuron input of the Hopfield network can be written as

$$\frac{du_i}{dt} = -\frac{\delta E}{\delta v_i} \quad (8)$$

or

$$\frac{du_i}{dt} = -\sum_{j=1}^N T_{ij} v_j + I_i. \quad (9)$$

The Hopfield network can therefore be used for energy minimization problems if the weights and biases are arranged such that they describe an energy function, with the minimum of energy occurring at the stable state of the network [20]. By specifying different values for the weights and biases, any hypothetical energy minimization problem can be simulated.

Many real world problems can be formulated as the minimization of an energy function, and this is central to the design of a Hopfield neural network formulated as an optimization tool. The energy function used must represent the problem correctly, and reach a minimum at the solution of the problem. Once this function is designed, the weights and biases can be set, and the network is built around these.

Most real world problems contain built-in constraints in addition to a goal that must be considered. These constraints form a cost added to the objective within the energy function, which can then be defined as

$$Energy = Goal + Constraints. \quad (10)$$

If the energy function is arranged in this particular way, the constraints become part of the minimization process, which means that the constraints do not need to be treated separately, just weighted by their importance to the problem. The Hopfield network process then finds the minimum energy that represents a compromise between the goal and the constraints.

The Hopfield network has been used within the field of remote sensing for ice mapping, cloud motion, and ocean current tracking [23], [24]. These demonstrate the utility of the Hopfield network for feature tracking, the basic principle of which

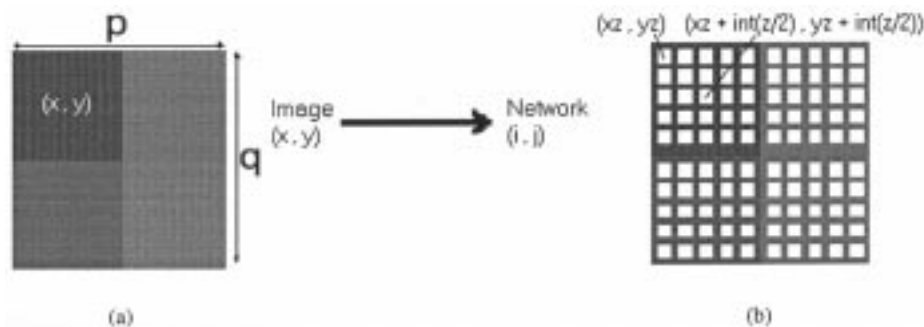


Fig. 3. (a) 2×2 pixel image, p and q represent the image dimensions, x and y represent the image pixel coordinates. (b) Representation of the Hopfield network for the image in (a). i and j represent the neuron coordinates ($\text{int} = \text{integer value}$).

is to match common features in a sequence of images. In addition, more general feature matching problems have been tackled using a Hopfield network [25]–[27], as well as applications such as recognition or classification [28], [29].

Hopfield and Tank [20] also used the network’s energy minimization capability and demonstrated solutions to complex combinatorial problems such as the traveling salesman problem. Hopfield and Tank formulated the problem as the minimization of an energy function, and for a ten-city problem, achieved convergence to a valid tour with an 80% success rate. Later, work in [30] and [31] eliminated the presence of local minima in the energy functions where constraints were not satisfied by setting various constants in a certain arrangement, and this guaranteed convergence to a valid solution.

III. USING THE HOPFIELD NEURAL NETWORK FOR TARGET IDENTIFICATION AT THE SUBPIXEL SCALE

The input data for the research described in this paper were derived from aerial photography, whereby targets were identified and extracted accurately from the photographs by hand, using field survey for verification. By degrading these verification images of clearly defined land cover targets to the spatial resolution of Landsat TM data using a square mean filter, accurate class proportion estimates were obtained for each pixel. These provided the input to the network, but in practice, this input could come from automated fuzzy classification methods such as the multilayer perceptron. However, for the research in this paper, the aim was to understand and test the capabilities of the Hopfield network technique, so any error introduced to the input data by an automated fuzzy classification method would be detrimental to this aim.

Mapping the spatial distribution of the class components within each pixel was formulated as a constraint satisfaction problem, and an optimal solution to this problem was determined by the minimum of an energy function coded into a Hopfield neural network. The network architecture was arranged to represent a finer spatial resolution image, and constraints within the energy function determined the spatial layout of binary neuron activations within this arrangement. The Hopfield neural network was used to find the minimum of the energy function, which corresponded to a bipolar map of class components within each pixel and this method is outlined in detail in the following sections.

A. Network Architecture

In many papers, on the use of Hopfield neural networks for optimization, the spatial relations between neurons are considered irrelevant. However, for this paper, the nature of the problem and the proposed solution requires the network neurons to be considered as being arranged in a regular grid, with positioning within this grid being of significance to the network design for this task (Fig. 3). Therefore, neurons will be referred to by coordinate notation, for example, neuron (i, j) refers to a neuron in row i and column j of the grid and has an input voltage of u_{ij} and an output voltage of v_{ij} . The zoom factor z determines the increase in spatial resolution from the original satellite sensor image to the new high-resolution image and after convergence to a stable state, the neurons represent a bipolar classification of the land cover target *at the higher spatial resolution*. Fig. 3 shows the notation used in this paper and how coordinates are transformed linearly from the image space to the network neuron space, for example, the pixel (x, y) in the satellite image is represented by $z \times z$ neurons centred at coordinates $[xz + \text{int}(z/2), yz + \text{int}(z/2)]$, where int is the integer value.

B. Network Initialization

Each neuron is initialized with a starting value u_{init} , and two strategies for initializing the network exist.

- 1) Each set of neurons representing a pixel in the low-resolution image is identified and a proportion of this set is randomly given an output of $u_{\text{init}} = 0.55$. This proportion is equal to the actual area proportion of the class within the image pixel and the remaining neurons of the set are given an output of $u_{\text{init}} = 0.45$. The values of 0.55 and 0.45 were chosen as the initial on and off outputs to speed up processing time and avoid unnecessary bias toward certain energy minimization paths. In [20] and many other papers related to the use of the Hopfield network for solving the traveling salesman problem, neurons are initialized with a random value close to the central state value (0.5). This choice is justified by the fact that no initial preference should be given to any path. The small difference between the two values also enables the network to “push” neuron outputs to 1 or 0 to represent a bipolar

classification faster than if, for example, a neuron was initially given an output of 0 and had to be pushed to 1 to produce an optimal solution.

- 2) The completely random initialization of neuron outputs within the range $u_{\text{init}} = [0.45, 0.55]$. This allows performance comparison with the class proportion-defined initialization, and does not introduce any possible unnecessary bias into the result, which may occur using 1), should estimated class proportions be inaccurate.

C. Implementation

When implemented on a digital computer, sets of biases and weights do not need to be determined, as the network is simulated via its equation of motion (9) using the Euler method

$$u_{ij}(t + dt) = u_{ij} + \frac{du_{ij}(t)}{dt} dt \quad (11)$$

where dt is the time step of the iterative method and the function $du_{ij}(t)/dt$ is measured using dE_{ij}/dv . Equation (8) shows the correspondence between the two functions, and dE_{ij}/dv is determined using the goals and constraint of the super-resolution target identification task. Equation (11) is run until $\sum_{i,j} u_{ij}(t + dt) - u_{ij}(t) \leq du_c$, where du_c is a sufficiently small value, and the equations of motion were defined as

$$\frac{dE_{ij}}{dv_{ij}} = k_1 \frac{dG1_{ij}}{dv_{ij}} + k_2 \frac{dG2_{ij}}{dv_{ij}} + k_3 \frac{dP_{ij}}{dv_{ij}}. \quad (12)$$

Each component of (12) is described in the subsequent sections.

D. The Energy Function

The goal and constraints of the subpixel mapping task were defined such that the network energy function was

$$E = - \sum_i \sum_j (k_1 G1_{ij} + k_2 G2_{ij} + k_3 P_{ij}) \quad (13)$$

where

k_1, k_2 and k_3 constants weighting the various energy parameters;

$G1_{ij}$ and $G2_{ij}$ output values for neuron (i, j) of the two objective (or goal) functions (see Section III-D1), and these correspond to the quadratic term in (5);

P_{ij} output value for neuron (i, j) of the proportion constraint (see Section III-D1), which corresponds to the linear term in (7).

1) *Goal Functions:* The goal (objective) functions were based upon an assumption of spatial order [18]. Almost all natural and human-made phenomena exhibit spatial continuity at some scale, such that points near to each other are more alike than those further apart, and the degree of dissimilarity depends on both the environment and the nature of our observations [32]. These observations can be the pixels in remotely sensed images, and the assumption of spatial order can be used to infer relationships between these pixels. By focusing within this research on discrete land cover targets, which all exhibit spatial order to some degree, the assumption of spatial order becomes

particularly relevant. Therefore, by devising simple measures of spatial order and incorporating each as objective functions within the Hopfield network, to map the spatial distribution of class components within a pixel, this real world phenomenon was modeled.

In this case, the aim of the goal was to make the output of a neuron similar to that of its neighboring neurons. Therefore, if the output of neuron (i, j) was similar to the average output of the eight neighboring neurons, then a low energy is given. If it is different, then this represents an undesirable situation in terms of the aim of spatial order, and a high energy is produced. However, to produce a bipolar image, a function that just drives a neuron output to be similar to the surrounding neuron output is insufficient. Consequently, two objective functions were introduced, one to increase neuron output toward a value of 1 and another to decrease neuron output to 0, each dependent on the average output of the eight neighboring neurons.

The first function aimed to increase the output of the centre neuron v_{ij} to 1 if the average output of the surrounding eight neurons $(1/8) \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl}$ was greater than 0.5

$$\frac{dG1_{ij}}{dv_{ij}} = \frac{1}{2} \left(1 + \tanh \left(\frac{1}{8} \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl} - 0.5 \right) \lambda \right) \times (v_{ij} - 1) \quad (14)$$

where λ is a gain which controls the steepness of the \tanh function. The \tanh function controls the effect of the neighboring neurons. If the averaged output of the neighboring neurons is less than 0.5, then (14) evaluates to 0, and the function has no effect on the energy function (13). If the averaged output is greater than 0.5, (14) evaluates to 1, and the $(v_{ij} - 1)$ function controls the magnitude of the negative gradient output, with only $v_{ij} = 1$ producing a zero gradient. A negative gradient is required to increase neuron output.

The second goal function aimed to decrease the output of the centre neuron $v_{ij} = 1$ to 0, given that the average output of the surrounding eight neurons $(1/8) \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl}$ was less than 0.5

$$\frac{dG2_{ij}}{dv_{ij}} = \frac{1}{2} \left(1 + \left(-\tanh \left(\frac{1}{8} \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl} \right) \lambda \right) \right) v_{ij}. \quad (15)$$

This time, the \tanh function evaluates to 0 if the averaged output of the neighboring neurons is more than 0.5. If it is less than 0.5, the function evaluates to 1 and the center neuron output v_{ij} controls the magnitude of the positive gradient output, with only $v_{ij} = 0$ producing a zero gradient. A positive gradient is required to decrease neuron output only when $v_{ij} = 1$ and $(1/8) \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl} > 0.5$, or $v_{ij} = 0$ and $(1/8) \sum_{\substack{k=i-1 \\ k \neq i}}^{i+1} \sum_{\substack{l=j-1 \\ l \neq j}}^{j+1} v_{kl} < 0.5$ is the energy gradient equal to zero, and $G1_{ij} + G2_{ij} = 0$. This satisfies the objective of

recreating spatial order, while also forcing neuron output to either 1 or 0 to produce a bipolar image.

2) *Proportion Constraint*: While the goal functions provide the enforcement of spatial order, the sole use of these functions would result in all neuron outputs taking the values 1 or 0. Therefore, a method of constraining the effect of those functions to the correct image areas was required. The proportion constraint P_{ij} aimed at retaining the pixel class proportions output from the fuzzy classification. This was achieved by adding in the constraint that the total output from the set of neurons representing each coarse resolution image pixel should be equal to the predicted class proportion for that pixel. An area proportion estimate representing the proportion of neurons with an output of 0.55 or higher was calculated for all the neurons representing pixel (x, y)

Area Proportion Estimate

$$= \frac{1}{2z^2} \sum_{k=xz}^{xz+z} \sum_{l=yz}^{yz+z} (1 + \tanh(v_{kl} - 0.55)\lambda). \quad (16)$$

The use of the \tanh function ensures that if a neuron output is above 0.55, it is counted as having an output of 1 within the estimation of class area per pixel. Below an output of 0.55, the neuron is not counted within the estimation, which simplifies the area proportion estimation procedure and ensures that neuron output must exceed the random initial assignment output of 0.55 in order to be counted within the calculations.

To ensure that the class proportions per pixel output from the fuzzy classification were maintained, the proportion target per pixel a_{xy} was subtracted from the area proportion estimate (16)

$$\frac{dP_{ij}}{dv_{ij}} = \frac{1}{2z^2} \sum_{k=xz}^{xz+z} \sum_{l=yz}^{yz+z} (1 + \tanh(v_{kl} - 0.55)\lambda) - a_{xy}. \quad (17)$$

If the area proportion estimate for pixel (x, y) is lower than the target area, a negative gradient is produced that corresponds to an increase in neuron output to counteract this problem. An overestimation of class area results in a positive gradient, producing a decrease in neuron output. Only when the area proportion estimate is identical to the target area proportion for each pixel does a zero gradient occur, corresponding to $P_{ij} = 0$ in the energy function (13).

E. Technique Advantages

The approach described in this paper holds several strategic advantages over those techniques mentioned in Section I-A.

- The option to choose the level of spatial resolution increase. This is essential if simulation of higher spatial resolution imagery is the aim.
- The Hopfield network technique has the ability to simulate any shape rather than being restricted to straight boundaries.
- The ease with which any additional information can be incorporated within the framework to aid the target identification. Any prior information about the land cover target depicted in the input imagery can be coded easily into the

Hopfield network as an extra constraint to increase accuracy.

- The design of the Hopfield network as an optimization tool means that all constraints are satisfied simultaneously, rather than employing a multistage operation.
- The effect that each one of these constraints has on the final prediction image can be controlled simply via weightings.
- While the incorporation of prior information on each land cover target may increase map accuracy, the Hopfield network technique has the benefit of being able to produce accurate, super-resolution land cover target maps from just class proportions. There is therefore no reliance on the availability of finer spatial resolution imagery or land line vector data.

IV. INTERPRETATION

To understand and illustrate the workings of the Hopfield network set up in this way, several synthetic images were created. Traditionally, within the remote sensing community, there has been a reluctance to use synthetic imagery, with the application of techniques directly to real imagery being preferred. However, by breaking down the elements of real-world imagery into simplified representations, understanding an image processing technique and, in turn, making improvements to it becomes easier.

The spatial order exhibited in all natural and human-made landscapes is demonstrated by the fact that scenes within remotely sensed imagery are composed of various shapes, for example, fields, roads and houses. The variety of the spatial order of shapes can be characterized using compactness and circularity. Compactness c is defined as

$$c = \frac{4\pi a}{p^2} \quad (18)$$

where p is the perimeter length, and a is the area of the shape.

Circularity r is defined as

$$r = \frac{a}{\pi(\max)^2} \quad (19)$$

where a is the area of the shape, and \max is the maximum distance from the shape centre to the perimeter. Fig. 4 shows eight, 56×56 pixel, synthetic images of differing shapes, with their respective compactness and circularity measures. The shapes represent extremes of spatial order as well as possible land cover targets, for example (b) a road and (d) a field. In addition, the shapes cover a wide range of compactness and circularity values, providing a useful test of the generalization capabilities of the Hopfield network. Fig. 5 shows each synthetic image subsampled (using a 7×7 mean filter) to generate an 8×8 pixel image, which causes mixing of the two classes (white and black) at the shape boundaries, producing eight proportion images and imitating the effect of class mixing within remotely sensed imagery. By using these eight proportion images as inputs to the Hopfield network and setting a zoom factor of 7, it should be possible to test the capabilities of the network by approximating the eight images each was derived from.




Fig. 4. Synthetic images and the features of the shapes depicted in each.

A. Setting the Constraint Weightings

To attempt to make predictions from the synthetic imagery, optimum values of goal and constraint weightings k_1 , k_2 and k_3 should be used, as these constants are of great importance because they control the direction of the optimization process. For this paper, equal weightings of 1.0 were chosen and the justification for this decision becomes clear with the following examples.

Fig. 6(a) shows a hypothetical situation of an image pixel with a zoom factor of five. It should be noted that for this example, the effect of proportions within surrounding pixels is ignored for simplicity, but in practice, the distribution of class proportions within the surrounding pixels may have a significant effect on the resulting neuron activations predicted.

The target class proportion is $8/25$, and using the constrained random initialization method, this is satisfied immediately (of the 25 neurons, eight are activated). While the proportion constraint value P is zero, the goal function values for this arrangement G_1 and G_2 are large due to the isolated outputs of neurons C and D. To minimize G_1 and G_2 and keep the target proportion of $8/25$, neurons A and B should have an increased output, while the outputs of C and D should be reduced.

If the proportion constraint is strongly weighted, then the arrangement of neural outputs displayed in Fig. 6(a) remains, as P is minimized, and the goal functions are not weighted strongly enough to have an effect, i.e., $k_1 G_1 + k_2 G_2 < k_3 P$.

If the goal functions are strongly weighted, then neurons A and B display increased output, neurons C and D stay at a high output, and other neurons around these also increase in output. In such a case, the proportion constraint is not sufficiently strong to maintain target class proportions, i.e., $k_1 G_1 + k_2 G_2 > k_3 P$.

The effect of such biased weightings is demonstrated by running the network with a zoom factor of 7, for 10 000 iterations on the synthetic image in Fig. 5(a). Fig. 6(b) demonstrates that the predicted shape is too large and irregular without the proportion constraint to control the positive activation effect of the goal functions. Giving P a large weight means that maintaining target class proportions becomes a priority, and the goal functions have little effect, such that the image in Fig. 6(c) is produced with a range of outputs.

By weighting the goal functions and proportion constraint equally, i.e., $k_1 = k_2 = k_3$, each affects and controls the other to minimize the overall energy of the network, and this is demonstrated in the following examples.

B. Predictive Ability and Limitations

The results of the Hopfield network predictions from the synthetic imagery after 10 000 iterations, using values of 1.0 for k_1 , k_2 and k_3 are shown in Fig. 7. Three measures of accuracy were calculated to assess the difference between each network prediction and the target images. One of the simplest measures of agreement between a set of known proportions \mathbf{y} , and a set of estimated proportions \mathbf{a} is the area error proportion (AEP) per class

$$AEP = \frac{\sum_{q=1}^n (y_q - a_q)}{\sum_{q=1}^n a_q} \quad (20)$$

where n is the total number of neurons. This statistic informs about bias in the prediction, and as it is based solely on area predictions, it represents a measure of the success of the proportion constraint in maintaining the target proportions.




Fig. 5. The 56×56 pixel synthetic imagery shown in Fig. 3 next to each degraded to 8×8 pixel images.

The correlation coefficient r represents an alternative measure of the amount of association between a target and estimated set of proportions

$$r = \frac{c_{y,a}}{s_y \cdot s_a}, \quad c_{y,a} = \frac{\sum_{q=1}^n (\bar{y}_q - y_q) \cdot (\hat{a}_q - a_q)}{n-1} \quad (21)$$

where $c_{y,a}$ is the covariance between y and a and s_y , and s_a are the standard deviations of y and a . This statistic informs about the variance of the subpixel spatial distribution.

Another standard measure is the root mean square error (RMSE) per class

$$RMSE = \sqrt{\frac{\sum_{q=1}^n (y_q - a_q)^2}{n}} \quad (22)$$

which informs about the accuracy of the prediction (bias and variance).

Fig. 7(a) demonstrates the predictive ability and generalization capabilities of the Hopfield network used in this paper. From the eight by eight pixel image in Fig. 4, the network is able to re-create perfectly the circle it was derived from Fig. 4(a). For a set area of pixels, the circle represents the maximum spatial order attainable, and so by basing the goal functions around this assumption, the network is able to perform well on such regular shapes.



Fig. 6. (a) Hypothetical image pixel with a zoom factor of five imposed. The pixel has been initialized with eight neurons of high output (white) and 17 with low output (black), (b) Hopfield network prediction for the shape in Fig. 4(a), given $k_1 = 1.0$, $k_2 = 1.0$ and $k_3 = 0.1$, (c) Hopfield network prediction for the shape in Fig. 4(a), given $k_1 = 0.1$, $k_2 = 0.1$, and $k_3 = 1.0$.

Fig. 7(b) demonstrates the limitations that exist when class proportions lie on the edge of images. The goal functions are set up to rely on information from surrounding neurons and the lack of such information for edge neurons, consequently means that only the proportion constraint is satisfied, leading to the poor predictive performance at the image edge.

Fig. 7(c) and (f) demonstrate good predictive ability in terms of class area, resulting in low area error proportion estimates. However, in both cases, the predicted shapes are slightly incorrect, resulting in a larger RMSE. This is due to certain edge formations satisfying the goal functions and proportion constraint, yet not being identical to the edges in the target image. The lack of information available to recreate such edges correctly shows that the problem is underconstrained.

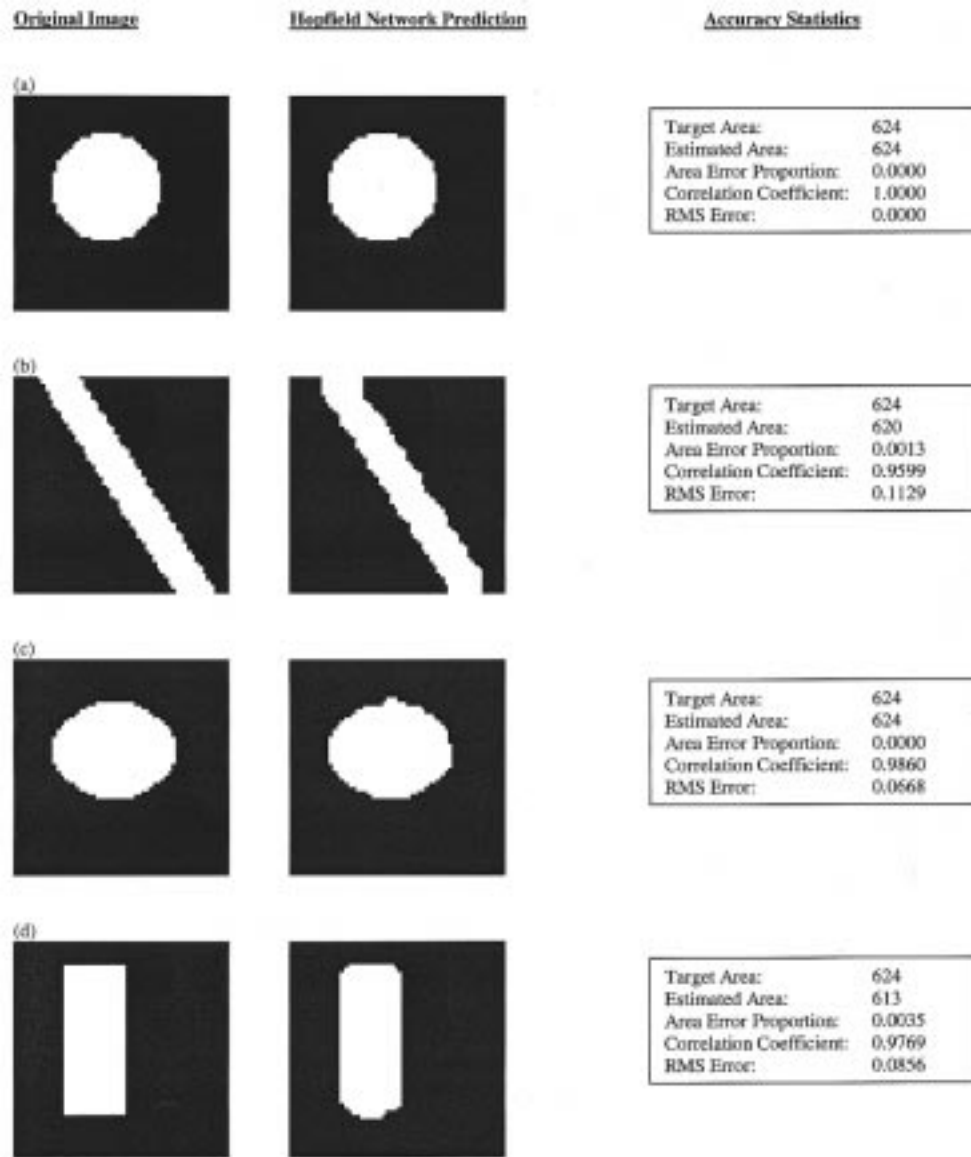


Fig. 7. Hopfield network predictions and accuracy assessment, given the 8×8 pixel synthetic imagery shown in Fig. 5 as input.

Fig. 7(d) reflects this problem further because, as described previously, the use of spatial order as the basis for the goal functions, means that the network will almost always converge to curved rather than sharp corners. Unless prior knowledge exists on the type of shapes that the network is aiming to recreate, this problem will remain.

Finally, Figs. 7(e), (g), and (h) demonstrate the ability of the network to cope with more complex shapes. Comparison of, for example, Figs. 7(e) with (f) shows identical area error proportion values of 0.0003, demonstrating how the class area has been predicted accurately. However, the correlation coefficients and RMS errors are significantly different, representing the difficulty in predicting the spatial distribution of class components for a complex shape such as that in Fig. 7(e). While the proportion constraint has ensured that class area is maintained, without prior information on the shape depicted in the input proportions, the goal functions on their own are insufficient to recreate accurately the spatial layout of the cross.

By repeating the network run on the more complex input proportions in Figs. 7(e), (g), and (h), slightly different predictions are produced each time. The use of an initialization technique based on random neuron output, constrained by target class proportions, means that starting neuron arrangements from certain network runs produce lower energy than others. This results in different paths of convergence along the energy surface of the network, and in turn, for the more complex shapes, this results in slightly different predictions each time, again indicating the underconstrained nature of the problem.

1) *Predictive Ability versus Shape Type:* Relationships can be drawn between the shape characteristics and the predictive ability of the network. These have the potential to be used to obtain a network performance prediction, providing that a certain degree of knowledge about the input land cover shapes is known.

Figs. 8(a) and (b) show the relationship between the various shape characteristics and the RMSE. The plots suggest

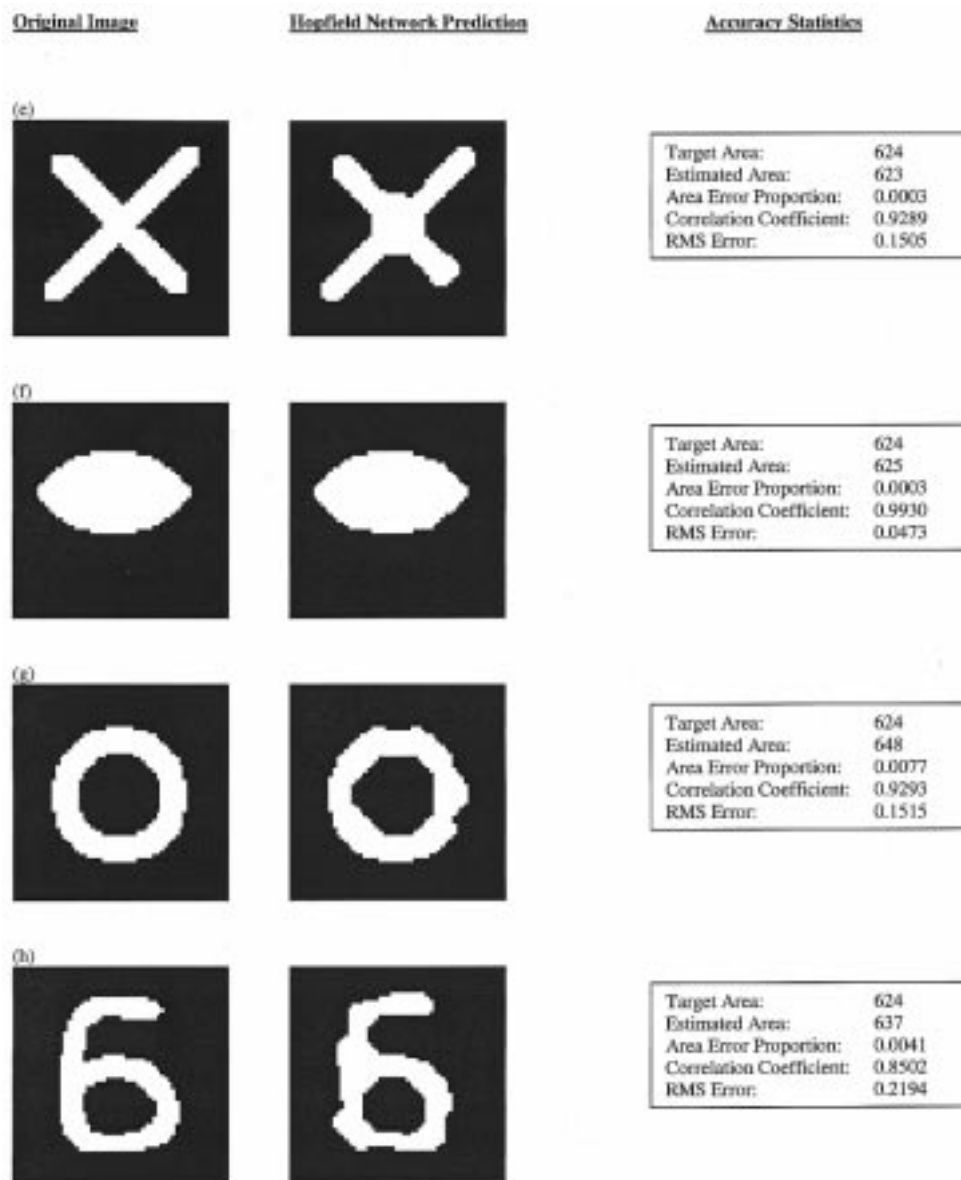


Fig. 7. (Continued.) Hopfield network predictions and accuracy assessment, given the 8×8 pixel synthetic imagery shown in Fig. 5 as input.

improved performance with compactness and circularity maximized, which represents maximum spatial order and is due to the design of the goal function, as mentioned previously.

2) *Predictive Ability versus Zoom Factor*: Fig. 9 demonstrates the effect that different zoom factors have on the ability of the Hopfield network to predict accurately the spatial layout of various shapes. The network was run on the synthetic imagery shown in Fig. 5, using zoom factors of 3, 5, 7, and 9, and the averaged results were plotted.

Fig. 9(a) shows little difference between the area error proportions for the different zoom factors, and this indicates that the zoom factor has little effect on the accuracy with which the areal coverage of the shape is predicted. However, Fig. 9(b) and (c) demonstrates the effect that the zoom factor has on the precision of the prediction, because as the zoom factor is increased, the accuracy with which the target shape is recreated increases.

The drawback is increased computational time and therefore, a balance must be found between the spatial resolution and accuracy required, and computational cost.

3) *Predictive Ability versus Iterations*: Fig. 10 shows the performance of the network over 10 000 iterations, after initialization using the proportion-constrained method. The network was run on the synthetic imagery shown in Fig. 5 using a zoom factor of 7 and values of 1.0 for k_1 , k_2 and k_3 , and the averaged results were plotted. The three plots demonstrate how, from the proportion-constrained random initialization, the major spatial organization of neuron outputs is undertaken within the first 1000 iterations. Fig. 10(a) shows how after 2000 iterations, the area error proportion reaches a stable value, reflecting the satisfaction of the proportion constraint, and this leaves the goal functions to be minimized, and reflected in the gradual convergence to stable values of Figs. 10(b) and (c).




Fig. 8. Scatterplots showing the variation in RMSE with (a) compactness and (b) circularity.

C. Initialization

The use of just three weighting functions in determining the energy of the Hopfield network described in this paper, enables the energy surface to be plotted. This simplicity means that the Hopfield network set up in this way is not only understandable, but also computationally efficient, and by studying the shape of such a surface, the workings of the network can be better understood. Fig. 11(a) shows energy plots for three different initializations of the network. The network was run on the 8 synthetic images shown in Fig. 5 using a zoom factor of 7, and values of 1.0 for k_1 , k_2 and k_3 , and the averaged results were plotted.

- i) Represents a typical energy path when the network is initialized using image class proportions. The proportion constraint is satisfied immediately, leaving the network to minimize the total energy by altering neuron outputs to reduce the overall goal value.
- ii) Represents a typical energy path when the neurons are initialized to random outputs, which are mostly less than 0.55, and the shape of the plot reflects the design of the proportion constraint. The threshold built into the proportion constraint means that, unless neuron output is greater than 0.55, it is not counted as “on” in terms of representing class proportions. The proportion constraint, therefore, has little effect until the goal functions increase the output above 0.55 of enough neurons. At this point, the goal functions are sufficiently satisfied that the network minimizes the total energy by altering neuron outputs to reduce the overall constraint function value. In satisfying the constraint function, the dominant energy-minimizing force shifts toward the goal functions,

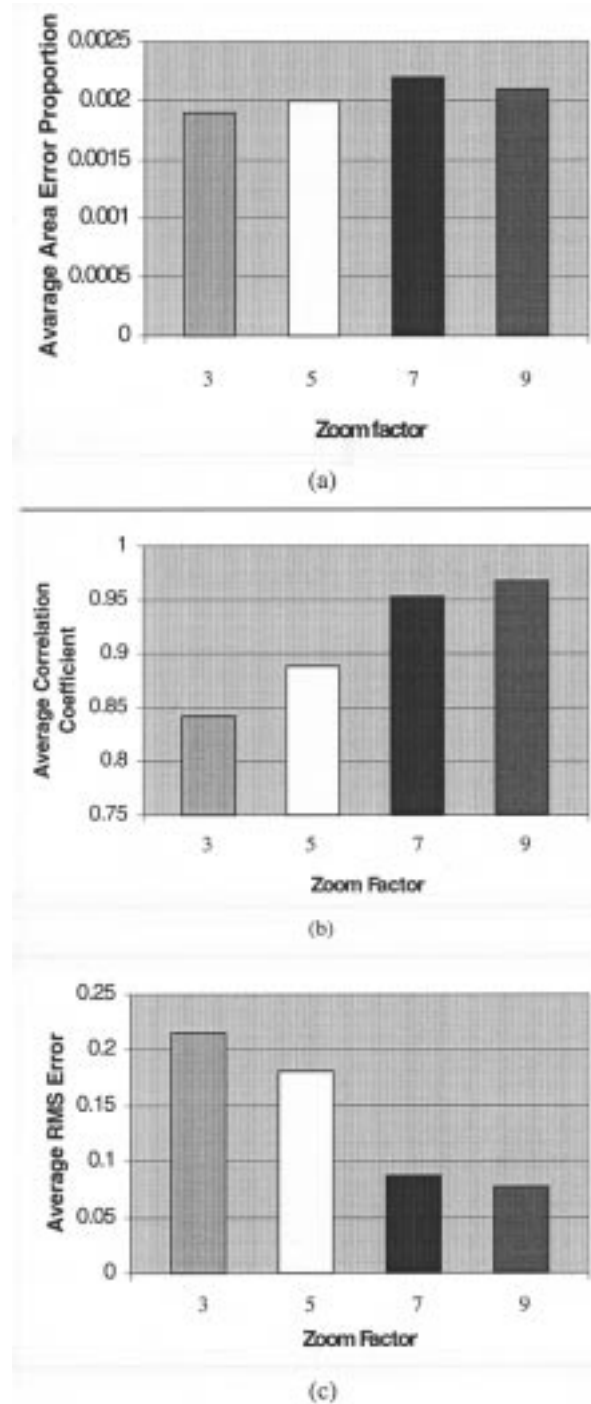


Fig. 9. Bar charts showing the relationship between zoom factor and (a) average area error proportion, (b) average correlation coefficient, and (c) average RMSE.

which again increase neuron output, until the network finds a balance between the two functions, and converges to an energy minimum.

- iii) Represents a typical energy path when the neurons are initialized to random values, which are mostly greater than 0.55. This satisfies neither the goal nor the constraint functions and the network, therefore, minimizes the total energy by altering neuron outputs to find a compromise between the functions.




Fig. 10. Graphs showing the relationship between number of iterations and (a) average area error proportion, (b) average correlation coefficient, and (c) average RMSE.

Fig. 11(b) displays a close-up of the convergence of the three energy plots in Fig. 11(a) and shows the three plots reaching approximately the same point, demonstrating the effectiveness of the network at converging to similar energy minima, given any initialization. Of the three runs depicted in Fig. 11(a), (i) reached the energy minimum quickest (approximately 5000 iterations), followed by (ii) (approximately 8000 iterations), then (iii) (approximately 9000 iterations).

D. Performance Predictions

This section has revealed several features about the workings of the Hopfield network that can be used to generate predictions on its performance.

- The technique will produce a more accurate prediction if the shape depicted in the input class proportions is compact and circular.
- The technique will produce a more accurate prediction if a high zoom factor is used.
- The technique will produce a more accurate prediction if the network is allowed to run for at least 1000 iterations.

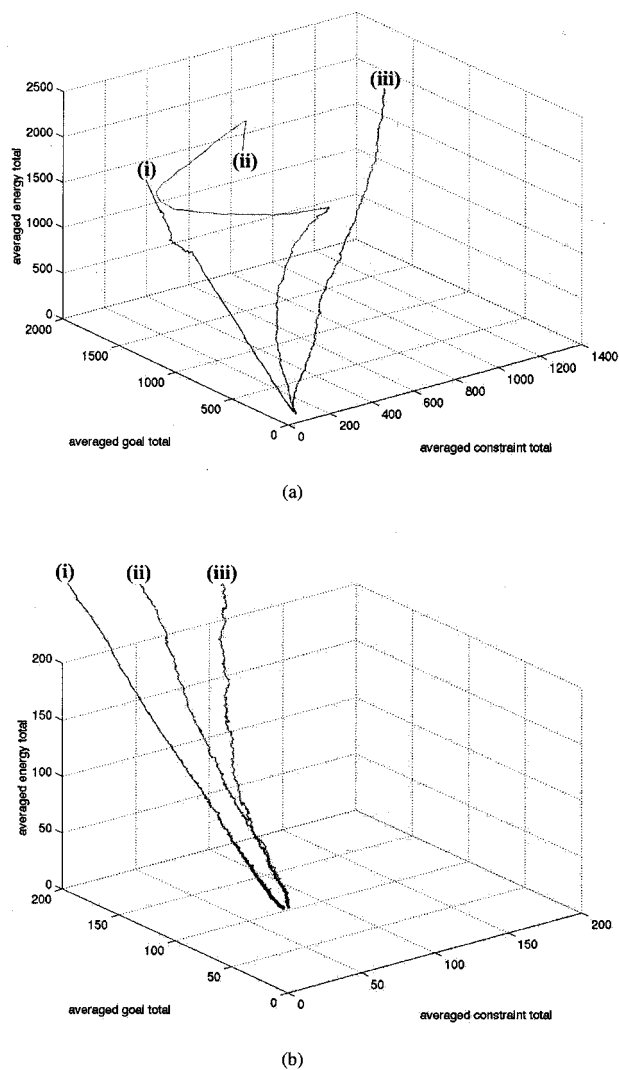


Fig. 11. (a) Hopfield network energy plots for three different initialization settings: (i), (ii), and (iii) and (b) close up of the convergence of the three plots shown in (a).

- The network will converge to an accurate prediction in fewer iterations if a proportion-constrained initialization is used.

V. RESULTS

Results were produced using the Hopfield network run on a P2-350 computer. The network was used to identify land cover targets at the subpixel scale from simulated remotely sensed imagery. Landsat TM imagery was acquired over an agricultural area east of Leicester (Stoughton), U.K., and seven wavebands at a spatial resolution of 30 m. Within the imagery, attention was focused on a wheat field and a section of airstrip, which provided clearly defined targets with which to evaluate the technique.

Figs. 12 and 13 show the various data used to initialize the network and evaluate the results produced. The verification data shown in Figs. 12(b) and 13(b) were derived by field survey and hand from the 0.5 m spatial resolution digital aerial photographs




Fig. 12. (a) Digital aerial photograph (1 km grid overlaid), (b) verification data derived from aerial photography and ground survey, (c) 24×24 pixel Landsat TM band 4 image, and (d) wheat class proportions.




Fig. 13. (a) Digital aerial photograph (1 km grid overlaid), (b) Verification data derived from aerial photography and ground survey, (c) 24×24 pixel Landsat TM band 4 image, and (d) asphalt class proportions.

shown in Figs. 12(a) and 13(a). The class proportion estimates shown in Figs. 12(d) and 13(d) were calculated from the verification data using a square mean filter that avoided the potential problems of incorporating error from the process of fuzzy classification of the imagery in Figs. 12(c) and 13(c).

The network was initialized using the wheat and asphalt proportion images shown in Figs. 12(d) and 13(d). The proportion-constrained initialization method was used with values for k_1 , k_2 and k_3 of 1.0, and zoom factors of 5 and 7 were used for comparison.

After 10 000 iterations of the network at zoom factor 5 (approximately 10 min running time), prediction images were produced [Figs. 14(a) and (c)] with spatial resolutions five times higher than that of the input class proportion images in Figs. 12(d) and 13(d). In addition, after 10 000 iterations of the network at zoom factor 7 (approximately 20 min running time), prediction images were produced [Figs. 14(b) and (d)] with spatial resolutions seven times finer than that of the input class proportion images in Figs. 12(d) and 13(d). The same measures of accuracy used in Section IV were calculated to assess the difference between each network prediction and the verification data.

VI. DISCUSSION

The high accuracies shown for the results in Fig. 14 indicate that the Hopfield network has the potential to locate accurately target class proportions within pixels from remotely sensed imagery.

The regularity and discrete nature of the wheat field enabled the network to perform well on this particular land cover target. In Fig. 14(a) and (b), the network maintained the areal coverage

of the field in its prediction, while accurately predicting the field shape also, and in both cases, the nature of the goal functions meant that the network predicted rounder corners than that of the actual field. With a zoom factor of seven, there was a less pronounced rounding effect due to the finer scale that the goal functions were working on, resulting in greater accuracy. This corner-rounding problem represents the under-constrained nature of the problem, and prior information about the field shape at its corners could potentially be built into the network as a constraint to avoid this problem.

The more complex shape of the airfield, as expected, produced results of lower accuracy than those for the wheat field. However, the statistics in Figs. 14(c) and (d) demonstrate that for both cases, the shape was predicted accurately, with an RMSE as low as 0.094 pixels using a zoom factor of seven. As predicted in Section IV-D, more accurate results were again produced using the higher zoom factor. Again, both predictions produced rounder corners than those of the actual land cover target, but by using a zoom factor of seven, the network was able to model more accurately these corners than when it was run with a zoom factor of 5.

VII. CONCLUSION

This study has shown that a Hopfield neural network can be used to estimate the location of the class proportions within pixels and produce a land cover target map of subpixel geometric precision. The Hopfield network represents a robust, efficient, and simple technique, and results from synthetic and simulated remotely sensed data show good performance, suggesting that it has the potential to identify accurately land cover targets at the subpixel scale.

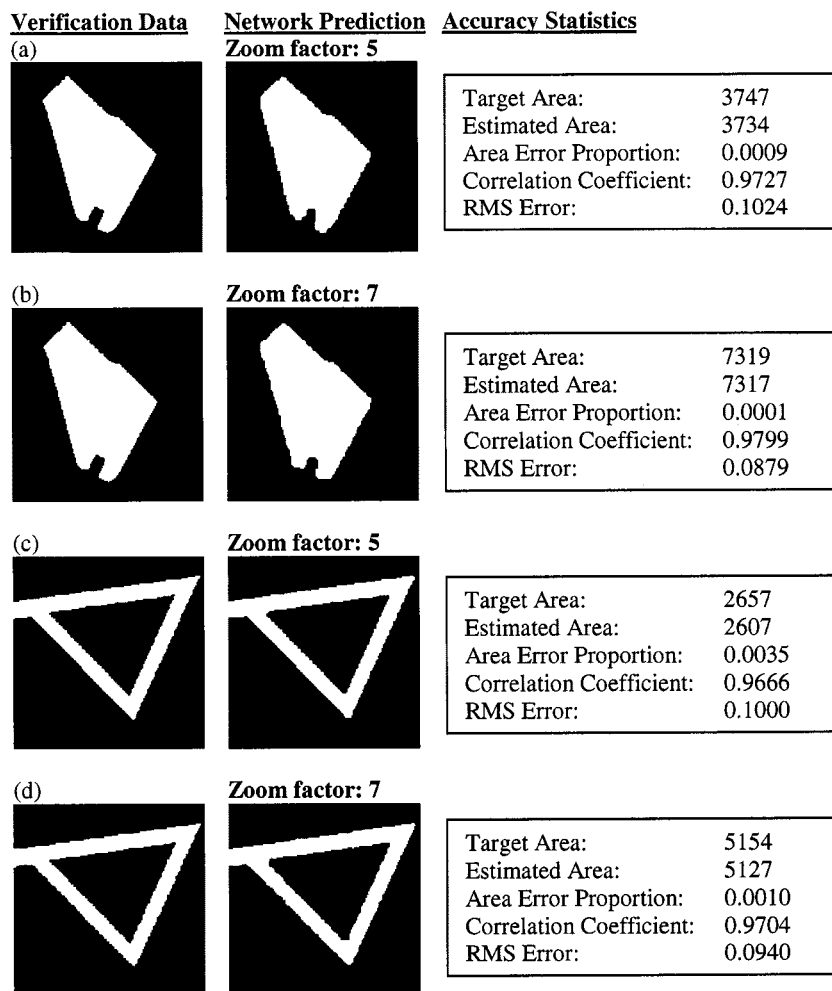


Fig. 14. Hopfield network predictions and accuracy assessment, given the class proportion images shown in Figs. 11 and 12 as input.

Further work could be undertaken, examining the performance of the technique on imagery of differing spatial resolutions to provide a measure of the applicability of the technique to imagery from sensors other than the Landsat TM. The findings from this paper suggest that the problems that do exist with the Hopfield network technique are due to the underconstrained nature of the problem, and additional future work may focus on incorporating various forms of prior information about the imagery to be processed as constraints in the network.

ACKNOWLEDGMENT

The authors would like to thank the E.U. FLIERS Project (ENV4-CT96-0305) for providing the imagery, and would also like to thank Drs. L. Bastin and M. Edwards, University of Leicester, Leicester, U.K.

REFERENCES

- [1] C. E. Woodcock and A. H. Strahler, "The factor of scale in remote sensing," *Remote Sens. Environ.*, vol. 21, pp. 311–332, 1987.
- [2] J. B. Campbell, *Introduction to Remote Sensing*, 2nd ed. New York: Taylor and Francis.
- [3] P. Fisher, "The pixel: A snare and a delusion," *Int. J. Remote Sensing*, vol. 18, pp. 679–685, 1997.
- [4] F. J. Garcia-Haro, M. A. Gilabert, and J. Melia, "Linear spectral mixture modeling to estimate vegetation amount from optical spectral data," *Int. J. Remote Sensing*, vol. 17, pp. 3373–3400, 1996.
- [5] P. M. Atkinson, M. E. J. Cutler, and H. Lewis, "Mapping subpixel proportional land cover with AVHRR imagery," *Int. J. Remote Sensing*, vol. 18, pp. 917–935, 1997.
- [6] R. A. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*. San Diego, CA: Academic, 1997.
- [7] M. Brown, S. R. Gunn, and H. G. Lewis, "Support vector machines for optimal classification and spectral unmixing," *Ecol. Modeling*, vol. 120, pp. 167–179, 1999.
- [8] W. Schneider, "Land use mapping with subpixel accuracy from landsat TM image data," in *Proc. 25th Int. Symp. Remote Sensing and Global Environmental Change*, Ann Arbor, MI, 1993, pp. 155–161.
- [9] J. Steinwendner and W. Schneider, "A neural net approach to spatial subpixel analysis in remote sensing," in *Proc. 21st Workshop of the Austrian Association for Pattern Recognition*, Halstatt, Austria, 1997.
- [10] —, "Algorithmic improvements in spatial subpixel analysis of remote sensing images," in *Proc. 22nd Workshop of the Austrian Association of Pattern Recognition*, Illmitz, Austria, 1998, pp. 205–213.
- [11] J. Steinwendner, W. Schneider, and F. Suppan, "Vector segmentation using spatial subpixel analysis for object extraction," *Int. Arch. Photogram. Remote Sensing*, vol. 32, pp. 265–271, 1998.
- [12] J. Steinwendner, "From satellite images to scene description using advanced image processing techniques," in *Proc. 25th Remote Sensing Soc. Conf.*, U.K., 1999, pp. 865–872.
- [13] W. Schneider, "Land cover mapping from optical satellite images employing subpixel segmentation and radiometric calibration," in *Machine Vision and Advanced Image Processing in Remote Sensing*, I. Kanellopoulos, G. Wilkinson, and T. Moons, Eds. Heidelberg, Germany: Springer-Verlag, 1999, pp. 229–237.

- [14] J. Flack, M. Gahegan, and G. West, "The use of subpixel measures to improve the classification of remotely sensed imagery of agricultural land," in *Proc. 7th Australasian Remote Sensing Conf.*, Melbourne, Australia, 1994, pp. 531–541.
- [15] V. F. Leavers, "Which Hough transform," *CVGIP—Image Understanding*, vol. 58, pp. 250–264, 1993.
- [16] P. Aplin, P. M. Atkinson, and P. J. Curran, "Fine spatial resolution simulated satellite imagery for land cover mapping in the United Kingdom," *Remote Sens. Environ.*, vol. 68, pp. 206–216, 1999.
- [17] P. M. Atkinson, "Mapping subpixel boundaries from remotely sensed images," in *Innovations in GIS IV*, Z. Kemp, Ed. London, U.K.: Taylor and Francis, 1997, pp. 166–180.
- [18] G. Matheron, *Les variables régionalisées et leur estimation*, Paris: Masson.
- [19] G. M. Foody, "Sharpening fuzzy classification output to refine the representation of sub-pixel land cover distribution," *Int. J. Remote Sensing*, vol. 19, pp. 2593–2599, 1998.
- [20] J. Hopfield and D. W. Tank, "Neural computation of decisions in optimization problems," *Biol. Cybern.*, vol. 52, pp. 141–152, 1985.
- [21] A. Cichocki and R. Unbehauen, *Neural Networks for Optimization and Signal Processing*. Stuttgart, Germany: Wiley.
- [22] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," in *Proc. National Academy of Sciences*, vol. 81, 1984, pp. 3088–3092.
- [23] S. Côté and A. R. L. Tatnall, "The Hopfield neural network as a tool for feature tracking and recognition from satellite sensor images," *Int. J. Remote Sensing*, vol. 18, pp. 871–885, 1997.
- [24] H. G. Lewis, "The Use of Shape, Appearance and the Dynamics of Clouds for Satellite Image Interpretation," Ph.D. Thesis, Univ. Southampton, Southampton, 1998.
- [25] N. M. Nasrabadi and C. Y. Choo, "Hopfield network for stereo vision correspondence," *IEEE Trans. Neural Networks*, vol. 3, pp. 5–13, Jan. 1992.
- [26] P. Forte and G. A. Jones, "Posing structural matching in remote sensing as an optimization problem," presented at the Machine Vision and Advanced Image Processing in Remote Sensing, I. Kanellopoulos, G. Wilkinson, and T. Moons, Eds., London, U.K., 1999.
- [27] R. Li, W. Wang, and H. Tseng, "Detection and location of objects from mobile mapping image sequences by Hopfield neural networks," *Photogramm. Eng. Remote Sensing*, vol. 65, pp. 1199–1205, 1999.
- [28] P. P. Raghu and B. Yegnanarayana, "Segmentation of Gabor-filtered textures using deterministic relaxation," *IEEE Trans. Image Processing*, vol. 5, pp. 1625–1636, 1996.
- [29] P. Campadelli, D. Medici, and R. Schettini, "Color image segmentation using Hopfield networks," *Image and Vision Computing*, vol. 15, pp. 161–166, 1997.
- [30] S. V. B. Aiyer, M. Niranjan, and F. Fallside, "A theoretical investigation into the performance of the Hopfield model," *IEEE Trans. Neural Networks*, vol. 1, pp. 204–215, 1990.
- [31] S. Abe, "Global convergence and suppression of spurious states of the Hopfield neural networks," *IEEE Trans. Circuits Syst.*, vol. 40, pp. 246–257, 1993.
- [32] P. J. Curran and P. M. Atkinson, "Geostatistics and remote sensing," *Progress Phys. Geogr.*, vol. 22, pp. 61–78, 1998.

Andrew J. Tatem received the B.Sc. degree from the University of Southampton, Southampton, U.K., in 1998. He is currently pursuing the Ph.D. degree in the Departments of Geography and Electronics, University of Southampton.

His research interests focus on land cover mapping at the subpixel scale from remotely sensed images.

Hugh G. Lewis received the B.Eng. (Hons.) and M.Sc. degrees in control engineering and control systems, respectively, in 1992 and 1993, from The University of Sheffield, Sheffield, U.K., and the Ph.D. degree from the University of Southampton, Southampton, U.K., in 1999 for remote sensing image analysis using neural networks.

Between 1997 and 1999, he investigated neural networks and empirical models for land area estimation from remote sensing data in the Department of Aeronautics and Astronautics, and the Department of Electronics and Computer Science, University of Southampton. In October 1999, he joined the School of Engineering Sciences to research the long-term evolution of the space debris environment of the Earth.

Peter M. Atkinson received the B.Sc. degree from the University of Nottingham, Nottingham, U.K., and the Ph.D. degree from the University of Sheffield, Sheffield, U.K., in 1986 and 1990, respectively.

He is currently a Reader in physical geography with the University of Southampton, Southampton, U.K. He teaches physical geography and GIS at undergraduate levels and has supervised eight Ph.D. students. Between 1989 and 1990, he worked at Logica U.K. Ltd., and from 1990 and 1993, he was a Postdoctoral Researcher with the University of Bristol, Bristol, funded by the Leverhulme Trust. He has been Principal Investigator or Co-Principal Investigator on several grants and contracts. He has published widely in refereed journals and has edited numerous books and special issues of journals on topics relating to geostatistics and remote sensing. His primary research interests are geostatistics, spatial statistics, and remote sensing.

Dr. Atkinson has sat on the Council of, and is a Fellow of, the Remote Sensing and Photography Society (RSPS) and founded the RSPS Models and Advanced Techniques Special Interest Group (MATSIG).

Mark S. Nixon is a Reader in the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K. He has helped to develop new techniques for static and moving shape extraction (both parametric and nonparametric), which have found application in automatic face and gait recognition, in medical image analysis, and in remotely sensed image analysis. He was the Principal Investigator for the University of Southampton's part of the EU Fuzzy Land Information from Environmental Remote Sensing (FLIERS) Programme. He has a new textbook on feature extraction in computer vision forthcoming with Butterworth Heinemann, U.K. His research interests include image processing and computer vision.

Dr. Nixon chaired the Ninth British Machine Vision Conference BMVC'98.