

Novelty detection and neural network validation

C.M. Bishop

Indexing terms: Novelty assessment, Neural networks

Abstract: One of the key factors which limits the use of neural networks in many industrial applications has been the difficulty of demonstrating that a trained network will continue to generate reliable outputs once it is in routine use. An important potential source of errors is novel input data; that is, input data which differ significantly from the data used to train the network. The author investigates the relationship between the degree of novelty of input data and the corresponding reliability of the outputs from the network. He describes a quantitative procedure for assessing novelty, and demonstrates its performance by using an application which involves monitoring oil flow in multiphase pipelines.

1 Introduction

Neural networks have been shown to have a useful degree of performance over a wide range of industrial and medical applications. However, a key factor which limits the widespread implementation of neural-network solutions in many areas has been the difficulty of demonstrating that the outputs generated by the network in the field are reliable. In general, the problem of network validation is a difficult one, and it involves many issues, some of which are generic to any software system. Here we consider only those aspects which are specific to neural networks and related methods.

It is useful to distinguish between two broad levels of network validation. At the first level we provide, in addition to the network outputs, some associated measure of confidence so that potentially unreliable outputs can be detected. In many applications this will give a useful degree of validation, since it could allow the bulk of the input data to be processed with high confidence, with remaining data either being discarded or processed by other means. An example would be the automation of medical image interpretation for mass screening programmes, in which images that gave network outputs with a low confidence could be rejected by the network and interpreted instead by human experts. The second level of validation requires a guarantee that the network outputs are reliable for any achievable input vector. This represents a much more rigorous degree of validation, and it would be appropriate for networks used in safety-critical applications, for example. In this paper we shall consider only the first level of validation.

© IEE, 1994

Paper 1330K (C4, E5), first received 29th November 1993 and in revised form 19th May 1994

The author is with the Neural Computing Research Group, Department of Computer Science, Aston University, Birmingham B4 7ET, United Kingdom

Intuitively we expect that a network will generate reliable results when presented with data which are similar to those used during training, but that when substantially novel data are presented the network outputs will be prone to serious error. In Section 2 we investigate the relationship between the novelty of input data and validity of network outputs, and we use this as the basis of a practical system for network validation [2]. The technique is illustrated in Section 3 using an example from the monitoring of multiphase flows in oil pipelines. Possible extensions of this approach are discussed in Section 4.

2 Network validation

Consider a feedforward network trained by minimising a sum-of-squares error function. If we denote the joint probability density functions for the training data by $p(\mathbf{x}, t_j)$, then we can write the error in the form

$$E = \sum_{j=1}^m \int [y_j(\mathbf{x}; \mathbf{w}) - t_j]^2 p(\mathbf{x}, t_j) d\mathbf{x} dt_j \quad (1)$$

where $j = 1, \dots, c$ labels the output units, \mathbf{x} is the input vector to the network, y_j denotes the output from unit j , and t_j is the target value for that unit. The network corresponds to a set of functional mappings $y_j(\mathbf{x}; \mathbf{w})$, parametrised by a set of weights and biases \mathbf{w} whose values are found by minimising E .

We note that the joint density $p(\mathbf{x}, t_j)$ can be factorised into the product of the unconditional density of the input data $p(\mathbf{x})$ and the conditional density of the target data $p(t_j | \mathbf{x})$. After some simple algebra, we can then rewrite the error of eqn. 1 in the form

$$E = \sum_{j=1}^m \int [y_j(\mathbf{x}; \mathbf{w}) - \langle t_j | \mathbf{x} \rangle]^2 p(\mathbf{x}) d\mathbf{x} + \sum_{j=1}^m \int \{ \langle t_j^2 | \mathbf{x} \rangle - \langle t_j | \mathbf{x} \rangle^2 \} p(\mathbf{x}) d\mathbf{x} \quad (2)$$

where we have defined the conditional averages of the target data as

$$\langle t_j | \mathbf{x} \rangle \equiv \int t_j p(t_j | \mathbf{x}) dt_j \quad (3)$$

$$\langle t_j^2 | \mathbf{x} \rangle \equiv \int t_j^2 p(t_j | \mathbf{x}) dt_j \quad (4)$$

Note that only the first term in eqn. 2 depends on the network weights. Provided that the functions $y_j(\mathbf{x}; \mathbf{w})$ have sufficient flexibility, for instance if they correspond to a network with a sufficient number of hidden units, then the minimum of this error function occurs when

$$y_j(\mathbf{x}; \mathbf{w}) = \langle t_j | \mathbf{x} \rangle \quad (5)$$

so that the network outputs represent the regression of the target data, conditioned on the input vector. Note

that we are implicitly assuming an infinitely large data set, as we have replaced sums over data points with integrals over distributions. For finite data sets, it is necessary to limit the complexity of the network to control the trade-off between the bias and the variance of the solution. In principle, the solution given by eqn. 5 would then be obtained by considering a sequence of ever larger data sets and correspondingly more flexible network functions, such that (in the limit) both the bias and the variance of the solution are made to vanish [3].

The neural network solution represented by eqn. 5 can be regarded as optimal in the sense that, if the training data were to be generated from a deterministic function with superimposed zero-mean noise, then the network would average over the noise and learn the underlying function. Similarly, for classification problems in which the training data are labelled using a 1-of- N target coding scheme, eqn. 5 implies that the network outputs represent the posterior probabilities of the input vector x belonging to the corresponding classes, and so they can again be regarded as optimal. Note that the residual value of the error function at the minimum is given by the second term in eqn. 2, which represents the average variance of the noise on the data. Since this term is independent of the network weights, it plays no role in network training.

The key point to note is that the first term in the error function of eqn. 2 is weighted by $p(x)$ which represents the unconditioned density of the input data. In a practical application the goal of network training is to find a good approximation to the regression by minimisation of a sum-of-squares error defined over a finite training set. From eqn. 2 we expect this approximation to be most accurate in regions of input space for which the density $p(x)$ is high, since only then does the error function penalise the network mapping if it differs from the regression. In regions of input space where $p(x)$ is small, there is little contribution to E if the network outputs $y_j(x; w)$ and the regression functions $\langle t_j | x \rangle$ differ significantly. A similar argument applies to other error functions, such as the cross-entropy.

This suggests that the unconditional density can provide an appropriate quantitative measure of novelty with respect to the training data set. If a new input vector falls in a region of input space for which the density $p(x)$ is high then the network is effectively interpolating between training data points, and the network performance will generally be good. If the input vector falls in a region of input space for which $p(x)$ is low, then the input is essentially novel and the network could easily generate erroneous outputs.

We therefore arrive at the following procedure for validating network outputs. The data which are used to train the network are also used to construct an estimate $\hat{p}(x)$ of the (unknown) density $p(x)$. Standard cross-validation (using partitions of the training set) may be used at this point to optimise the network topology and the values of regularisation parameters, and this can also serve to optimise any smoothing parameters in the density model $\hat{p}(x)$. When the network is in use, each new input is presented to the trained network, and is also used to evaluate $\hat{p}(x)$ to provide a quantitative measure of the degree of novelty of the input vector. Input vectors x which have relatively small values of $\hat{p}(x)$ are those which are likely to generate spurious outputs. This is illustrated schematically in Fig. 1.

From eqn. 2 we see that the sum-squared error between the network outputs and the regression is

weighted by a variance factor of the form

$$\sigma_y(x) = \{p(x)\}^{-1/2} \quad (6)$$




Fig. 1 Use of the unconditional probability density to measure the degree of novelty of new input vectors

x denotes the vector of input variables to the network, the circles denote the training data, and the curve depicts a model $\hat{p}(x)$ for the unconditional density estimated from the training data points. A new input A , for which $\hat{p}(x)$ is large, is considered to be similar to the training data, while an input B , for which $\hat{p}(x)$ is small, is considered novel

and this can be used to assign x -dependent error bars to the network outputs $y_j(x)$, via the model $\hat{p}(x)$. Thus for each new input vector x we obtain the values of the network outputs y_j , together with a value σ_j for the confidence associated with this value. Note that this confidence value reflects the uncertainty in the value of the network estimate for the regression, and it is distinct from the variance of the target values around their conditional mean as determined by the intrinsic noise on the target data.

In some applications it will be more appropriate to place a threshold on $\hat{p}(x)$ and to reject all new data points for which $\hat{p}(x)$ falls below this threshold. The justification for such an approach is as follows. In applying a threshold we are implicitly classifying all new input vectors into one of two classes: those which are similar to the training data, which we denote by class C_1 ; and those which are novel, which we denote by class C_2 . The training data are assumed to be drawn entirely from C_1 , whereas new data could arise from either class, with prior probabilities $P(C_1)$ and $P(C_2)$, respectively, where $P(C_1) + P(C_2) = 1$. Given a new input vector x , we wish to assign it to one of the two classes in such a way as to minimise the probability of misclassification. This is achieved when the vector is assigned to the class with the larger posterior probability [4], so that the vector is assigned to class C_1 if

$$P(C_1 | x) > P(C_2 | x) \quad (7)$$

and to class C_2 otherwise. From the Bayes theorem the posterior probability of the vector x to belong to class C_i is given by

$$P(C_i | x) = \frac{p(x | C_i)P(C_i)}{\pi(x)} \quad (8)$$

where the unconditional density $\pi(x)$ is given by

$$\pi(x) = P(x | C_1)P(C_1) + P(x | C_2)P(C_2) \quad (9)$$

which acts as a normalisation factor and ensures that the posterior probabilities sum to unity.

$$P(C_1 | x) + P(C_2 | x) = 1 \quad (10)$$

Since the denominator in eqn. 8 is the same for both classes, it can be omitted from the comparison of probabilities needed to assign x to a particular class. Using

eqns. 7 and 8, we therefore assign x to class C_1 if

$$p(x|C_1) > \frac{p(x|C_2)P(C_2)}{P(C_1)} \quad (11)$$

and otherwise to class C_2 . The quantity $p(x|C_1)$ is the probability density from which the training data were drawn, which we have modelled by $\hat{p}(x)$. The density $p(x|C_2)$ represents the distribution of novel data, and so by definition we can have little idea of what form it should take. The simplest assumption is to take it to be constant over some large region of input space, falling to zero outside this region to ensure that the density function can be normalised. (In many applications, such as the one discussed in the next Section, the range of possible input values will in fact be bounded.) In this case the condition 11 is then equivalent to a threshold on $\hat{p}(x)$. This is illustrated in Fig. 2.




Fig. 2 Bayesian formalism for determining whether a new input vector is novel

It is assumed that the training data have arisen from class C_1 with prior probability $P(C_1)$ and distribution $p(x|C_1)$ (which we model using $\hat{p}(x)$), and that novel data are drawn from class C_2 with prior probability $P(C_2)$ whose distribution function $p(x|C_2)$ is taken to be constant over some large region of input space. From Bayes theorem, the optimal decision boundary for classifying new data as novel is equivalent to a threshold on $p(x|C_1)$. This divides the input space into two (possibly disjoint) regions R_1 and R_2 such that new input vectors which fall in the region R_2 are classified as novel

More generally, we might suppose that the density of novel data should be smaller in regions of input space where the density $\hat{p}(x)$ is large. Thus we could choose $p(x|C_2)$ to be of the form $F(p(x|C_1))$, where $F(\cdot)$ is a monotonically decreasing function. It is easily seen that this too gives rise to a threshold criterion on $p(x|C_1)$, which takes the form

$$p(x|C_1) > G^{-1}\left(\frac{P(C_1)}{P(C_2)}\right) = \text{const.} \quad (12)$$

where $G(z) \equiv F(z)/z$. Since $F(\cdot)$ is monotonic, so too is $G(\cdot)$, and hence G^{-1} is unique.

In a practical application of neural networks, an independent test set (one not used in the training and cross-validation process) is generally used to confirm the performance of the trained network. The same test can also be used to find a suitable threshold value for $\hat{p}(x)$. This is achieved by evaluating $\hat{p}(x)$ for all points in the test set, and then choosing a value for the threshold such that most (or all) of these points are classified as belonging to C_1 . This procedure will be illustrated in the next Section.

Many conventional techniques exist for constructing probability density estimates $\hat{p}(x)$ from finite samples [4, 9] and various adaptive 'neural' approaches have also been suggested. In the next Section we illustrate these ideas using one of the standard methods for density estimation.

3 Example application

To illustrate the concept of novelty detection as the basis of a practical system for network validation, we consider a specific industrial application of neural networks concerned with the determination of the oil fraction in a multiphase oil pipeline. Full details of this application can be found in Reference 1.

To reduce costs, the oil industry makes use of multiphase pipelines to transfer mixtures of oil, water and gas directly from offshore production fields to an onshore facility where the various components can be separated. This has led to the need for an effective noninvasive method for monitoring the oil fraction in such pipes, with sufficient accuracy for reservoir monitoring and custody transfer purposes. The problem is a complex one owing to the fact that the multiphase flow can exhibit a wide variety of configurations, and that numerical modelling of such flows is notoriously difficult.

One technique to determine the phase fractions in the pipeline uses multiple-beam dual-energy gamma densitometry [1], which is based on the attenuation properties of gamma beams passing through the pipe. The degree of attenuation of a gamma beam depends on the particular material in the path of the beam (expressed in terms of an absorption coefficient), the length of material through which the beam passes, and the gamma energy (i.e. wavelength). Fig. 3 shows a schematic illustration of

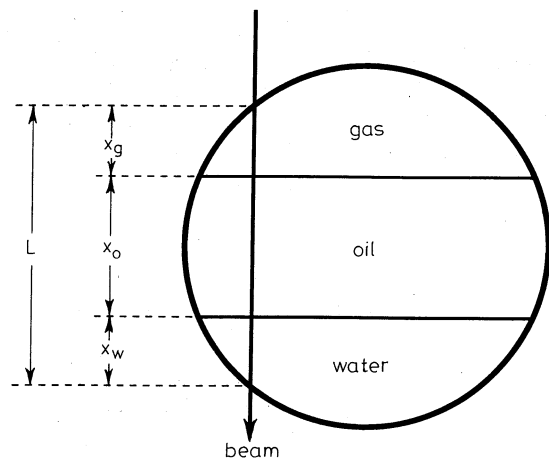


Fig. 3 Gamma beam passing through a pipe containing water, oil and gas in a stratified configuration

a collimated gamma beam passing through a circular cross-section pipe containing water, oil and gas in a stratified configuration.

We wish to determine the path lengths x_w , x_o and x_g in each of the three phases. By measuring the attenuation of the beam at two energies we obtain two independent pieces of information. The constraint that the three path lengths must sum to the total path length, so that $x_w + x_o + x_g = L$, provides a third piece of information. Knowing the absorption coefficients of the gamma beam in each of the three phases at each energy, it is then straightforward to evaluate the required path lengths [1].

If measurements could be made along many transverse paths through the pipe, it would in principle be possible to perform a tomographic reconstruction of the phase configuration, and hence to calculate the oil fraction. In a practical system, however, only a few lines of sight will be available, and so alternative analysis methods need to be used.

We have investigated linear mapping methods, as well as neural networks based on the multilayer perceptron

(MLP), to analyse the outputs from the densitometer [1]. It is found that the nonlinearity offered by the MLP gives a useful improvement in performance over the linear approach, and reduces the RMS error on new data by a factor of about two. We have considered a system with three vertical and three horizontal beams, with each of the six detectors generating two signals which correspond to the attenuation at each of the two wavelengths. These are first preprocessed to extract the fractional path lengths along the lines of sight in the oil and water phases. (The fractional path lengths in the gas phase are not considered because they represent redundant information). The resulting 12 numbers are used as inputs to a multilayer perceptron with a single hidden layer of sigmoidal units followed by two linear output units whose activations represent the fractions of oil and water in the pipeline.

For the purposes of this study synthetic data have been generated using configurations selected at random from the four examples shown in Fig. 4. Note that these

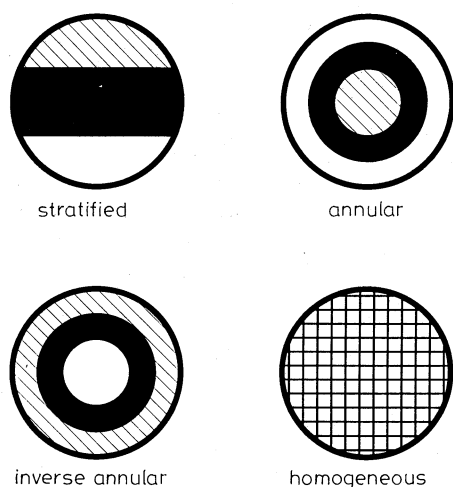


Fig. 4 Cross-sections of a pipe illustrating the four multiphase configurations used for network training

Data generated from a fifth configuration, 'inverted-stratified', are used to represent novel data

- ▨ gas
- oil
- water

configurations are chosen mainly for computational simplicity, and they do not necessarily give accurate representations of real multiphase flows. They are, however, sufficiently representative for the purposes of the present study. An important feature of this synthetic data is that the distribution of noise can be accurately modelled, as discussed below. To generate a data point, a configuration is chosen at random with equal probability from the four shown in Fig. 4. The oil, water and gas fractions are then also selected at random by choosing three numbers F_o , F_w and F_g at random with uniform distributions in the range (0, 1), and then calculating the oil fraction f_o using

$$f_o = \frac{F_o}{F_o + F_w + F_g} \quad (13)$$

with similar expressions for the water and gas fractions. This ensures that the three fractions sum to unity. The path lengths are then calculated, allowing for the effects of noise. In a practical system, the dominant contribution to the noise arises from the photon statistics, owing to the limited integration time of the detectors (necessary to

avoid averaging over temporal variations in the flow) and the limited source strength of the gamma beams, and has been accurately modelled with the correct Poisson statistics.

For a training set with 1000 points, the best generalisation performance is obtained from a network with eight hidden units. Fig. 5 shows a plot of the predicted oil




Fig. 5 Scatter plot of the oil fraction predicted by the neural network against the actual oil fraction

The original test set contains 1000 points

fraction against the actual oil fraction for 1000 points in an independent test set. It is clear that the network can predict the oil fraction with high accuracy. In a practical application of this technique the network would be trained using data from a multiphase flow test rig in which a variety of conditions, typical of those which arise in practice, can be produced.

To determine the degree of novelty of the input data we need a procedure to estimate the unconditional probability density $p(x)$ which corresponds to the training data set. For simplicity we have chosen a standard Parzen window approach with Gaussian kernel functions [9]. We shall discuss alternative approaches in Section 4. The Parzen estimator has one kernel function for each point in the training set, with its centre given by the corresponding input vector. For Gaussian kernels this gives a density model of the form

$$\hat{p}(x) = \frac{1}{n(2\pi)^{d/2}\sigma^d} \sum_{q=1}^n \exp \left\{ -\frac{|x - x^q|^2}{2\sigma^2} \right\} \quad (14)$$

where x^q represents a data point from the training set (which consists of n points in total), and d is the dimensionality of input space (so that here $d = 12$). The smoothing parameter σ controls the degree of smoothness of the estimated density function, and its value must be neither too large (since this gives the estimate a large bias) nor too small (since this gives the estimate a large variance). We have adopted the simple heuristic of setting σ to the average distance of the ten nearest neighbours from each point in the training set, averaged over the whole training set. This gave a value of σ close to that obtained using cross-validation.

To test the performance of this novelty detector we have generated a further data set which consists of 1000 examples with randomly chosen oil and water fractions, each corresponding to a fifth 'inverted-stratified' configuration which is obtained by inverting the stratified configuration of Fig. 4.

The value of the density function $\hat{p}(x)$ for a new value of x is generally called a likelihood. Fig. 6 shows a plot of

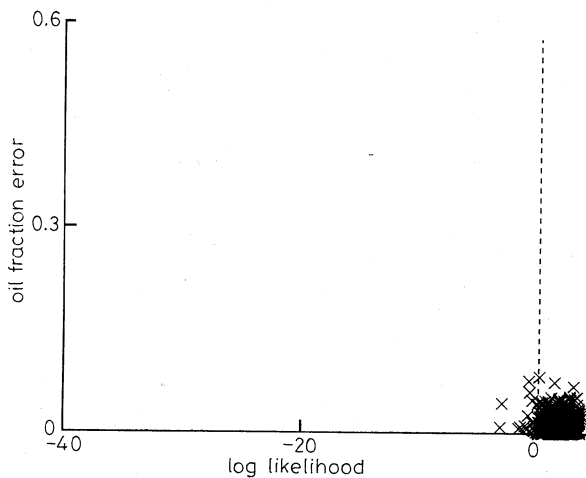


Fig. 6 Magnitude of the oil fraction error from the neural network prediction against the log likelihood from the novelty detector
1000 points from the original test set are shown

the log likelihood against the magnitude of the error between the oil fraction predicted by the neural network and the true value obtained from the data set, for the 1000 points from the test set used to plot Fig. 5. We see that the points all have large values for the log likelihood, and that the network output errors are correspondingly small. By comparison, Fig. 7 shows the corresponding

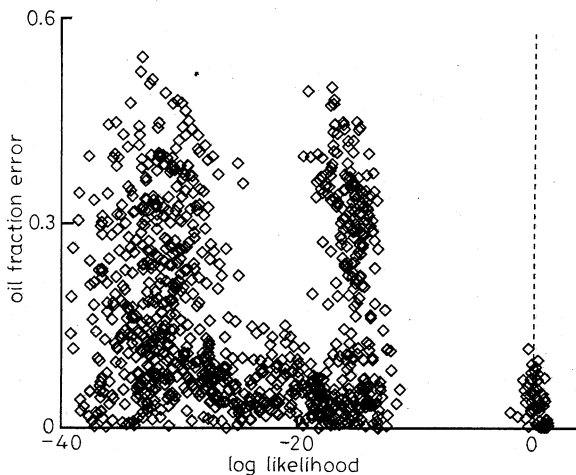


Fig. 7 Magnitude of the oil fraction error from the neural network prediction against the log likelihood from the novelty detector
1000 points corresponding to inverted-stratified configurations are shown

plot (on the same scale) for the 1000 examples of the novel inverted-stratified configuration. The majority of these points have log likelihood values which are substantially smaller than those of the test set points, and a corresponding larger range of oil fraction errors. We see that the network can indeed generate poor results when presented with data from this new configuration. Such data points can, however, easily be rejected on the basis

of their log likelihood values. Setting a threshold anywhere between -5 and -10 would reject all data points with significant phase fraction errors.

It can also be seen from Fig. 7 that there are some inverted-stratified points which have log likelihood values comparable to those from the original test set. Examination of the phase fractions for these points shows that they represent configurations which could also be classified as stratified configurations. For instance, if the oil phase fraction is sufficiently large then the three horizontal beam lines pass through oil only, and there exist stratified and inverted-stratified configurations which have the same phase fractions and which give rise to the same 12 path length measurements. The novelty detector 'correctly' interprets these as being similar to the training data, and indeed the network predicts the phase fraction with high accuracy.

4 Discussion

In this paper we have considered the problem of validating the outputs generated by a trained neural network, once it has been installed in an application. One of the most important sources of high errors in network predictions arises from novel input data, i.e. data which differ significantly from those used to train the network. We have shown that novelty can be quantified by modelling the unconditional probability density of the input data used during training, and by subsequently evaluating this density for all new data points. This can be used to assign error bars to the network outputs, or it can be used to classify input vectors on the basis of a novelty threshold.

We have illustrated the technique of novelty detection using a simple kernel-based estimator of the unconditional input density. Such an estimator is easy to implement and fast to train, but it is computationally expensive to evaluate for new data points, since the number of kernel functions is equal to the number of data points in the training set. An alternative approach is to use a semiparametric estimator, constructed from a Gaussian mixture model of the form

$$\hat{p}(x) = \sum_{i=1}^m \alpha_i \phi_i(x) \quad (15)$$

where

$$\phi_i(x) = \frac{1}{(2\pi)^{d/2} |\Sigma_i|^{1/2}} \exp \{ -(x - \mu_i)^T \Sigma_i^{-1} (x - \mu_i) \} \quad (16)$$

where Σ_i is the covariance matrix, and the mixing coefficients α_i satisfy

$$\sum_{i=1}^m \alpha_i = 1 \quad (17)$$

The parameters α_i , μ_i and Σ_i are adaptive parameters of the model, and their values are chosen by maximising the log likelihood of the training data with respect to the model $\hat{p}(x)$, given by

$$\log \mathcal{L} = \sum_{q=1}^n \log \hat{p}(x^q) \quad (18)$$

This can be done by using standard optimisation algorithms (such as conjugate gradients) or by using re-estimation techniques based on the EM algorithm [6]. The optimum value for the number of components m can be found by cross-validation. By comparison with the

simple Parzen method used in Section 3, the Gaussian mixture model requires significantly more computational effort to train. However, since the value of m is generally much smaller than the number of training data points, the computation time needed to evaluate $p(x)$ will typically be much less than with the Parzen estimator. Techniques for density estimation in which the number of basis functions is grown incrementally as part of the learning process, have been proposed in Reference 10, and these can also be used to measure novelty. In this case an appropriate threshold to detect novelty arises naturally as part of the training process.

We have discussed the validation of network outputs in terms of the detection of novelty with respect to the network input vectors. However, many practical applications of neural networks require that the raw input data be preprocessed before presentation to the network. This is especially important for problems with many input variables, from which a relatively small number of features are usually extracted. In such cases the novelty detection should be performed at the inputs to the pre-processing stage. This is important because the pre-processing itself is generally not information-preserving (for instance, it will often involve a reduction in dimensionality). Thus an input vector which is novel could be mapped to a vector at the inputs to the network which is similar to the training data, and which therefore would not be detected as novel.

An exception to this arises, however, when the pre-processing is constructed on the basis of some form of prior knowledge which is believed to hold accurately. For example, we know that the classification of the image of a character should be invariant to translations. In such cases, novel inputs whose novelty arises solely from application of the transformation should be accepted, because we expect the network to generate a satisfactory response, and so the novelty detection should be applied after the preprocessing.

The above discussion also makes it clear why (in general) it is incorrect to measure novelty in the space of activations of the hidden units in a network, since the transformation from inputs to hidden units is not in general information-preserving. There is one situation, however, in which the information provided by the activations of the hidden units can be used in the determination of novelty. The hidden units of a radial basis function network can be configured using unsupervised techniques based on density estimation in the input space [5, 8]. This might, for instance, be done using the approach discussed above to adapt the parameters of a Gaussian mixture model. In this case the resulting basis function $\phi_i(x)$ can be used directly as the hidden units in a radial basis function neural network, in which the second-layer weights are found by minimising a sum-of-squares error function using linear techniques. If the values of the mixing coefficients α_i , found from the optimisation of the basis functions, are retained they can be used to provide an estimate of the likelihood of new input data. In this case the α_i can be considered as weights to an extra output unit whose activation, given by eqn. 15, represents the likelihood. Such a network

combines in a natural way the generation of predictions, and the validation of those predictions, into a single network structure, as indicated in Fig. 8. (A related




Fig. 8 Radial basis function network in which the basis functions are used both in the evaluation of the network outputs $y_k(x)$, and to provide an estimate $\hat{p}(x)$ of the likelihood of input data for validation of the outputs

scheme, called the 'validity index' network, was proposed in Reference 7. An important difference is that the technique described above makes use of a principled method for density estimation based on maximum likelihood with respect to a Gaussian mixture model, rather than the more complex heuristic approach discussed in Reference 7.) Note, however, that the optimum choice of basis functions for density estimation need not be the optimum for prediction of the values of y_k .

5 References

- 1 BISHOP, C.M., and JAMES, G.D.: 'Monitoring of multi-phase flows using dual-energy gamma densitometry and neural networks', *Nucl. Instrum. Methods Phys. Res.*, 1993, **327**, pp. 580-593
- 2 BISHOP, C.M.: 'Neural network validation: an illustration from the monitoring of multi-phase flows'. Proceedings of the IEE Third International Conference on *Artificial Neural Networks*, Brighton, UK, 1993, pp. 41-45
- 3 BISHOP, C.M.: 'Neural networks and statistical pattern recognition' (Oxford University Press, Oxford, 1995, to appear)
- 4 DUDA, R.O., and HART, P.E.: 'Pattern classification and scene analysis' (John Wiley and Sons, New York, 1973)
- 5 LOWE, D.: 'On the iterative inversion of RBF networks: a statistical interpretation'. Proceedings of the IEE Second International Conference on *Artificial Neural Networks*, Bournemouth, UK, 1991, pp. 29-33
- 6 DEMPSTER, A.P., LAIRD, M.N., and RUBIN, D.B.: 'Maximum likelihood from incomplete data via the EM algorithm', *J. Royal Statist. Soc.*, 1977, **39**, pp. 1-38
- 7 LEONARD, J.A., KRAMER, M.A., and UNGAR, L.G.: 'Using radial basis functions to approximate a function and its error bounds', *IEEE Trans. Neural Networks*, 1992, **3**, pp. 624-627
- 8 NOWLAN, S.J.: 'Maximum likelihood competitive learning', in TOURETZKY, D.S. (Ed.): 'Advances in neural information processing systems' (Morgan Kaufmann, 1990), pp. 574-582
- 9 SILVERMAN, B.W.: 'Density estimation' (Chapman and Hall, New York, 1986)
- 10 ROBERTS, S., and TARASSENKO, L.: 'A probabilistic resource allocating network for novelty detection', *Neural Comput.*, 1993, **6**, pp. 270-284