

Number 475



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

A comparison of state-of-the-art classification techniques with application to cytogenetics

Boaz Lerner, Neil D. Lawrence

October 1999

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<https://www.cl.cam.ac.uk/>

© 1999 Boaz Lerner, Neil D. Lawrence

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<https://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

Abstract

Several state-of-the-art techniques: a neural network, Bayesian neural network, support vector machine and naive Bayesian classifier are experimentally evaluated in discriminating fluorescence in-situ hybridization (FISH) signals. Highly-accurate classification of signals from real data and artifacts of two cytogenetic probes (colours) is required for detecting abnormalities in the data. More than 3,100 FISH signals are classified by the techniques into colour and as real or artifact with accuracies of around 98% and 88%, respectively. The results of the comparison also show a trade-off between simplicity represented by the naive Bayesian classifier and high classification performance represented by the other techniques.

1 Introduction

In recent years, fluorescence in-situ hybridization (FISH) has emerged as one of the most significant new developments in the analysis of human chromosomes. FISH offers numerous advantages compared with conventional cytogenetic techniques since it allows numerical chromosome abnormalities to be detected during normal cell interphase. One of the most important applications of FISH is dot counting, i.e., the enumeration of signals (also

called dots) within the nuclei. Dot counting is used for studying numerical chromosomal aberrations in e.g., haematopoietic neoplasia, various solid tumours, prenatal diagnosis and for demonstrating disease-related chromosomal translocations (1).

However, a major limitation of the FISH technique for dot counting is the need to examine large numbers of cells. This is required for an accurate estimation of the distribution of chromosomes over cell population, especially in applications involving a relatively low frequency of abnormal cells. As visual evaluation by a trained cytogeneticist of large numbers of cells and enumeration of hybridization signals is expensive and time-consuming, FISH analysis for dot counting can be expedited by using an automatic procedure (2).

One approach to dot counting relies on an auto-focusing microscope to select the 'clearest' image for the analysis (2). However, basing dot counting on auto-focusing has some shortcomings (3). Instead, it has been recently proposed (3) to base FISH dot counting on images that are sampled at a fixed focal plane. The method enables most of the shortcomings of auto-focusing to be overcome, since it shortens the length of image acquisition and requires no special instrumentation. However, since the system captures images that contain many more unfocused signals, its ability to distinguish between focused ('real') signals and unfocused ('artifact') signals needs to be better than

that of a system employing an auto-focusing mechanism. Therefore, the proposed system depends upon two components: well-discriminating features to represent valid and artifact signal data and a highly-accurate framework to classify the signals.

Our previous work (4) has investigated the first component of feature representations for FISH signals. In the present work, we experimentally compare state-of-the-art classifiers in discriminating valid and artifact signals of two colours. We divide the classification procedure into two: classification of signals into colour and classification of signals as ‘real’ and ‘artifact’. In both cases, two-class classifiers are sought.

Section 2 of the paper describes stages of FISH image analysis that precede signal classification. Section 3 presents the four two-class trainable classifiers: a neural network, Bayesian neural network, support vector machine and naive Bayesian classifier that are evaluated for FISH signal classification. The experimental study is depicted in Section 4, while conclusions for this work are summarised in Section 5.

2 FISH image analysis

2.1 FISH acquisition

The process of preparing, hybridizing and screening FISH samples, as well as the procedure of capturing FISH images were described in Lerner *et al.* (3). A total of 400 images were collected from five slides and stored in TIFF format. An example of a FISH image used for dot counting is shown in Figure 1, where red and green fluorophores (signals), corresponding to chromosomes 21 and 13 respectively, are seen on blue stained nuclei.

2.2 FISH colour image processing and segmentation

By analysing each of the three colour channels – red, green and blue (RGB) of a FISH image separately and in various combinations, image processing can be facilitated. Nuclei can be analysed using the blue channel of the RGB image, whereas red and green signals are analysed separately in the red and green channels, respectively.

Segmentation on each of the three channels of the RGB image using global thresholds yields the image nuclei and red and green signals (4). Noise elimination and boundary smoothing of nuclei, as well as spatial correlation between nuclei and signals, complete the segmentation.

2.3 Signal feature measurement

Following segmentation, signals are characterized by sets of pixel intensities. A set (signal) can include one or many members (contiguous pixels). Since the content and dimension of each set can vary dramatically from signal to signal, raw data (intensities) are not considered discriminating enough to act as features for classification. It is therefore necessary to determine a more discriminating and compact representation of the data. One representation can be derived by measuring a set of features of the signal. The features include area (a size measure) and eccentricity (a shape measure), which have been previously suggested (2). In addition, we measure a number of spectral features (4). We compute, at the specific colour plane, three RGB intensity-based measurements: the total and average channel intensities and the channel intensity standard deviation. Following the conversion of RGB to HSI colour format (4), we can also compute four HSI (hue, saturation, intensity) based measurements: maximum hue, average hue, hue standard deviation, and delta hue. Delta hue is the difference between the maximum and average hue normalized by the average hue. This feature has been added to the set because it was observed that the difference between values of the average and maximum hue for real signals is usually near zero, whereas for some kinds of artifacts (e.g. overlap of two different fluorophores) this difference is sub-

stantially large. Two additional features of the set are the two coordinates of the eigenvector corresponding to the largest eigenvalue of the red and green intensity components of the signal. The last feature is the signal average grey intensity. More details about these features and a motivation for their selection are given elsewhere (4). Table 1 lists the twelve features used in the work.

3 An Overview of Several State-of-the-Art Two-Class Classifiers

Consider a training dataset D which consists of N data points with binary class labels $\{t_1 \dots t_N\}$ and vectors of inputs $\{\mathbf{x}_1 \dots \mathbf{x}_N\}$. We assume that the data was generated by some true underlying function $y(\mathbf{x})$. Our objective is to learn the parameters θ of some approximating function $f(\mathbf{x}, \theta)$ whose form is dependent on our model choice, \mathcal{M} , so that we may make ‘good predictions’ about our class labels.

3.1 Neural networks

For a neural network (NN), we choose a model which predicts the posterior probability of class membership. We define a likelihood function as (5)

$$p(D|\boldsymbol{\theta}) = \prod_{n=1}^N f(\mathbf{x}_n, \boldsymbol{\theta})^{t_n} [1 - f(\mathbf{x}_n, \boldsymbol{\theta})]^{(1-t_n)}. \quad (1)$$

The approximating function, f , is represented by the output of an NN with H hidden nodes in its single hidden layer

$$f(\mathbf{x}_n, \boldsymbol{\theta}) = \sigma \left(\sum_{h=1}^H v_h g(\mathbf{u}_h^T \mathbf{x}_n) \right). \quad (2)$$

The parameters $\boldsymbol{\theta}$ have been split into the input to hidden weights represented by H vectors \mathbf{u}_h , each vector being the weights that ‘fan-in’ to hidden node h , and \mathbf{v} , the vector of the hidden to output weights, consisting of H elements v_h . We have omitted biases for notational simplicity. We take the activation function g to be a hyperbolic tangent, and $\sigma(\cdot)$ is the logistic sigmoid function

$$\sigma(z) = \frac{1}{1 + \exp(-z)} \quad (3)$$

which constrains the output of the network to be between 0 and 1 allowing us to interpret f as the probability $P(C_1|\mathbf{x})$ that an input vector \mathbf{x} belongs to class C_1 .

We may now define an ‘error function’ as the negative log likelihood

leading to the cross entropy error function

$$-\ln p(D|\boldsymbol{\theta}) = -\sum_{n=1}^N \{t_n \ln(f(\mathbf{x}_n, \boldsymbol{\theta})) + (1 - t_n) \ln(1 - f(\mathbf{x}_n, \boldsymbol{\theta}))\}. \quad (4)$$

This error function may be minimised by a gradient-based optimisation method.

By using enough hidden units we may obtain a training error of zero. However, the resulting network will not generalise well on previously unseen data. We need to resort to a sub-partition of the training data known as a validation set to determine an appropriate number of hidden units. The unseen data are then tested on a network with that number of hidden units.

3.2 Bayesian neural networks

In the previous section we reviewed maximum likelihood learning for NNs. Now we briefly introduce the Bayesian approach for inferring parameters of an NN. In Bayesian learning we take our parameters to be random variables and we aim to determine their *posterior* distribution¹ for use in making predictions

$$p(\boldsymbol{\theta}|D) = \frac{p(D|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int p(D|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}}, \quad (5)$$

¹This posterior distribution should not be confused with the posterior probability of class membership.

where the likelihood $p(D|\boldsymbol{\theta})$ has already been defined in Equation 1, and we must specify a *prior* distribution $p(\boldsymbol{\theta})$ for the weights. The use of the prior is the main area of controversy in the Bayesian approach since it requires an interpretation of probabilities as being equivalent to ‘beliefs’ (6). The classical frequency-based definition of probabilities does not allow for such an interpretation. The prior represents our belief about what the weights should be before we see the data. A zero mean Gaussian prior is often used

$$p(\boldsymbol{\theta}) = \mathcal{N}(\mathbf{0}, \frac{1}{\alpha} \mathbf{I}) \quad (6)$$

where α is a ‘hyper-parameter’ which specifies the inverse width of the prior. This prior reflects a belief that negative weights are as likely as positive weights, and that smaller weights are more probable than larger weights, leading to smoother functions and better generalisation. Once the prior has been selected the integration in Equation 5 may be performed. To make predictions we look at the expected output of the network under the posterior distribution

$$\langle f(\mathbf{x}, \boldsymbol{\theta}) \rangle = \int f(\mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta}|D) d\boldsymbol{\theta}, \quad (7)$$

which is the posterior probability of C_1 for a new input vector \mathbf{x} .

Unfortunately for neural networks, the integrals required are non-analytic to compute and we must resort to approximations. There are various ap-

approaches, in particular the Laplace approximation (7; 8), variational approaches (9; 10; 11) and Monte-Carlo sampling (12). We follow the latter approach which involves obtaining samples $\{\boldsymbol{\theta}_1 \dots \boldsymbol{\theta}_S\}$ from the posterior distribution, and using them to make sample-based approximations to the required expectations,

$$\langle f(\mathbf{x}, \boldsymbol{\theta}) \rangle \approx \frac{1}{S} \sum_{s=1}^S f(\mathbf{x}, \boldsymbol{\theta}_s). \quad (8)$$

In particular, we can employ hybrid-Monte-Carlo techniques which make use of gradient information in the sampling.

3.3 The support vector machine

The support vector machine is a technique for classification and regression which arises from the field of *statistical learning theory*. In this work we only give a brief overview of the support vector machine (for a more detailed introduction see (13; 14)). We define a ‘loss function’, L , which is the penalty for a classification mistake. Consider first the *risk functional*, which is the expected value of the loss function under the joint probability of the data

$$R(\boldsymbol{\theta}) = \int L(t, f(\mathbf{x}, \boldsymbol{\theta})) p(t, \mathbf{x}) dt d\mathbf{x} \quad (9)$$

Ideally, to obtain the best approximation $f(\mathbf{x}, \boldsymbol{\theta})$ we would minimise the risk functional with respect to $\boldsymbol{\theta}$. Unfortunately the probability distribution

of the data $p(t, \mathbf{x})$ is unknown. Instead, we can look to a sample-based approximation to Equation 9 known as the empirical risk,

$$R_{\text{emp}}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^N L(t_n, f(\mathbf{x}_n, \boldsymbol{\theta})), \quad (10)$$

where we have assumed that our observed data points are independently drawn from the same distribution. The empirical risk may be minimised to obtain an approximation to $y(\mathbf{x})$. This is equivalent to the maximum likelihood approach for neural networks discussed in section 3.1.

Statistical learning theory relies on notions of capacity which reflect the number of patterns that a classifier may store. If the capacity of our classifier is not infinite, the value of the empirical risk will converge to that of the risk functional as the number of data points tends to infinity (15). Furthermore, it is possible to place bounds on the rate of convergence which hold with a certain confidence, normally taken to be 95%. These bounds are functions of the model capacity which is often quantified in terms of the *VC dimension* (16)

$$R(\boldsymbol{\theta}) \leq R_{\text{emp}}(\boldsymbol{\theta}) + R_{\text{struct}}(\mathcal{M}) \quad (11)$$

where R_{struct} is known as the structural risk and is a function which increases with increasing model capacity.

The principal of structural risk minimisation is to minimise not only the

empirical risk but also the structural risk, through capacity reduction, to obtain the best classifier. This is the underlying theory of the support vector machine.

In the case of binary classification, the support vector machine aims to place a separating hyper-plane between the two classes. Naturally, there are cases where a linear decision boundary does not exist and for these cases we must look to *kernel functions*. The kernel functions allow us to project the data onto a very high-dimensional feature space in which it may be separable by a hyper-plane. This has the side effect of also giving the model a very high capacity which increases the structural risk. Fortunately we may control this capacity by increasing the margin, thereby decreasing the structural risk. As increasing the margin also decreases the capacity, we select the solution with the maximum margin (see Figure 2). We may increase the margin further by allowing errors in the training set. If a training error is made the classifier pays a penalty which is proportional to the extent of the error. The constant of proportionality is often denoted by C . Different solutions will then be found depending on our choice of C (Figure 3).

A common kernel choice is a radial basis function. This kernel is based upon a Gaussian and requires the specification of a width parameter σ . The width parameter and the error penalty, C , may be set through the use of a

validation set.

3.4 The naive Bayesian classifier

For problems where the task is to minimise the probability of misclassification, the naive Bayesian classifier (NBC) provides a simple and clear method, while still enabling impressive performance. The NBC is termed naive since it makes use of a simplifying assumption that its observable variables, which represent the pattern features, are conditionally independent given the class variable. The classifier can be viewed as a special form of a Bayesian network (17), in which all the edges are directed from the class variable to the observable variables (Figure 4).

The NBC consists of a finite set $U = \{X_1, X_2, \dots, X_m, C\} = \{\mathbf{X}, C\}$ of random variables, where X_1, \dots, X_m are the observable variables that represent the features, and C is the class variable with K states. The NBC assigns a test pattern \mathbf{x} to the class C_k ($k = 1, \dots, K$) with the highest *posterior* probability

$$P(C_k|\mathbf{x}) = \frac{p(\mathbf{x}|C_k)P(C_k)}{p(\mathbf{x})} \propto p(\mathbf{X} = \mathbf{x}|C_k)P(C_k) = \prod_{i=1}^m p(X_i = x_i|C_k)P(C_k) \quad (12)$$

where $p(\mathbf{x}|C_k)$ is the *class-conditional* probability density, $P(C_k)$ is the *prior* probability for class C_k , $p(\mathbf{x})$, the *unconditional* density, normalises the pos-

terior probability such that $\sum_k P(C_k|\mathbf{x}) = 1$, $\mathbf{X} = \mathbf{x}$ represents the event that $X_1 = x_1 \wedge X_2 = x_2 \wedge \dots \wedge X_m = x_m$ and $\prod_{i=1}^m p(X_i = x_i|C_k)$ is the *likelihood* for \mathbf{x} . To derive this equation, we have omitted $p(\mathbf{x})$ which is common to all the states of the class variable and used the NBC independence assumption.

Both $P(C_k)$ and $p(\mathbf{x}|C_k)$ can be estimated from the data. $P(C_k)$ is the relative frequency of patterns belonging to C_k out of all the patterns in the data. The data can also be used to estimate $P(x|C_k)$ – the one-dimensional class-conditional probability for discrete variables and $p(x|C_k)$ – the one-dimensional class-conditional probability density for continuous variables. Modelling of probabilities is given by the sample frequency for each value of the variable (that is the number of times the value is observed divided by the total number of observations). Densities are estimated using different techniques, for example, single density estimation (a parametric method), kernel density estimation (a non-parametric method) or a Gaussian mixture model (a semi-parametric method) (5).

It has been found (18) that kernel density estimation (KDE) of the class-conditional probability densities of the FISH data is more accurate than the other two estimation methods. We model densities for each class C_k and observable variable X_m using a finite number of data points \mathbf{x}^n , $n = 1, \dots, N_k$,

where N_k is the number of training patterns in class C_k . KDE models the one-dimensional class-conditional density as a linear combination of kernel (usually Gaussian) functions

$$p(x|C_k) = \frac{1}{N_k} \sum_{n=1}^{N_k} \frac{1}{(2\pi h^2)^{1/2}} \exp \left\{ -\frac{\|x - x^n\|^2}{2h^2} \right\} \quad (13)$$

with width h centred around each of the training data points x^n of class C_k .

KDE, as other non-parametric methods, models non-normal distributed data more accurately than parametric techniques but at the cost of storage and computational complexities, as the number of variables in the model grows linearly with the number of training data points.

4 The experimental study

Before beginning the experiments, we established a database of 400 FISH images, which were captured from five slides. Following nucleus and signal segmentation, 3,144 objects were identified as potential signals and features were measured for them. Based on labels provided by expert inspection (see below), 1,145 of the signals were considered as ‘reals’ (among them 551 were red) and 1,999 as ‘artifacts’ (among them 1,224 were red).

Experiments to compare the accuracy of the four classification techniques on signals represented by the twelve features of Section 2.3 are conducted.

The normalised ($N(0, 1)$) signal features are classified in the first experiment into their colour – red or green, and in the second experiment as real or artifact. Therefore, the input and output spaces in both the experiments are twelve and one-dimensional, respectively. Labels for the patterns, as belonging to each of the classes, are needed to train and evaluate the classifiers, and they are obtained by an expert cytogeneticist using a custom-built graphical environment for labelling FISH images (19).

The NN classifier is a two-layer perceptron trained by the scaled conjugate gradient algorithm (5). Classification is based on the approximation of the multi-layer perceptron (MLP) outputs to the posterior probabilities for the classes. The two networks, colour NN and real NN used in the two experiments differ in the number of hidden units (H). To achieve the highest generalisation capability, this number is determined using a validation set. This set also insures that no over-training is performed during the 200 epochs of the training session. Finally, the classifier accuracy is averaged for each network over three random initialisations.

For the Bayesian neural network (BNN), the weights were split into four groups: the input to hidden weights, the hidden to output weights, the hidden layer biases and the output layer bias. A hyper-parameter α was associated with each group. For the output layer bias the hyper-parameter was fixed at

1×10^{-4} . The hyper-parameters for the other three groups were treated in a Bayesian manner. A gamma prior was specified across α :

$$p(\alpha) = \frac{a^b \alpha^{a-1} \exp(-b\alpha)}{\Gamma(a)}. \quad (14)$$

For the input to hidden layer weights, the hidden layer biases and the hidden to output layer biases the parameter a was taken to be 0.25 and the parameter b was taken to be 6.25×10^{-4} , 6.25×10^{-4} and 3.9×10^{-9} respectively. These values provide very broad hyper-priors, and were used in both the colour classification and the real-artifact classification. Two hundred samples were drawn from the posterior distribution using hybrid Monte-Carlo sampling (we used the implementation of Neal (20)). The last 150 samples were used for making predications on the test set. Since a prediction for a test pattern is the integral (average) network prediction over all the models considered under the prior (Equation 8), there is no need for validation; any necessary model selection takes place through the Bayesian framework.

For the support vector machine, two parameters needed to be set using a validation set. We considered models with an error penalty, C , of 1, 10, 100 and 1000. Using the radial basis function kernel, we explored widths, σ , of 5, 10, 50, 100, 500 and 1000 (we used the *SVM^{light}* implementation of Joachims (21)). These ranges of parameters provide a satisfactory environment to test the FISH data.

By defining a two-class classification task and setting the feature set to include all the twelve features, we also determine the NBC structure. If we select the class and observable variables of the NBC to represent these classes and features, respectively, we can employ the NBC for FISH signal classification. Then, we only need to estimate the class-conditional probability densities for each variable given each of the two states of the class variable for both the colour NBC and real NBC. For the variable that represents the area, which is the only discrete feature in the set, we model the class-conditional probability using the feature sample frequency. Class-conditional probability densities are modelled using KDE. Different width parameters (h) of the Gaussian kernels that are equal to $T/\sqrt{N_k}$ (where N_k is the number of training data points in C_k and T is in the range 0.01 to 10,000) are checked on the validation set. This choice guarantees that the parameter shrinks to zero as the number of instances goes to infinity and density estimation becomes increasingly local as the number of training points increases (22).

The experiments to evaluate the classification accuracy are conducted using the hold out method. This method is applicable for large data sets like the FISH data, which contains more than 3,100 data points. We partition the data randomly into training and test sets in proportion 70%/30%, respectively. We then select a model – H , h or C and σ for each of the clas-

sifiers – NN, NBC or SVM, respectively, using a validation set drawn from the training set and a five-fold cross-validation (CV-5) experiment. That is, for each model evaluated, we divide the training set into five disjoint sets, use four sets for training the classifier and the remaining set for validation. We repeat this procedure five times using all the different possible validation sets and average the classification accuracy of the five experiments for that model. The procedure is iterated for all the examined models. The model that obtains the highest average classification accuracy on the validation set is selected. Then, we re-train the classifier using the selected model and all the training data points and evaluate it on the test set. This procedure is repeated using the same partitions of the data for all the classifiers (except for the BNN that needs no validation). Figures 5 and 6 present results of model selection experiments performed with the NN and NBC, respectively. The classification accuracy on the training and validation sets is plotted for different models of the two techniques when the patterns are classified into their colour (by colour NN or colour NBC) and as ‘reals’ and ‘artifacts’ (by real NN or real NBC).

Finally, Table 2 compares the accuracy of each of the four techniques, using its own selected model, in classifying test FISH signals. Signals are discriminated by a classifier of colour (into red or green) and a classifier into

real or artifact. The comparison reveals that the BNN is the most accurate technique in both cases, and the NN and SVM are comparable and second best. The inferiority of the NBC compared with the other techniques is attributed to the relatively large amount of dependency among features of the set (e.g, average and maximum hue, total and average channel intensities). This dependency violates the independence assumption of the NBC and thereby decreases the classifier accuracy. However, due to their higher complexities, the other techniques can extract additional discriminating information from these features.

5 Conclusions

Highly-accurate signal classification is required for precise dot counting in FISH images that are captured without an auto-focusing microscope (3). To cope with this requirement, we apply in this paper state-of-the-art classification techniques – a neural network, Bayesian neural network, support vector machine and naive Bayesian classifier to the FISH data.

The four trainable classifiers discriminate data based on different approaches. NN training is based on maximum likelihood which is equivalent to the minimisation of an error function. Since the lowest training error does

not necessarily represent the model that provides the best generalisation, we are required to draw a validation set from the training set in order to perform model selection. This complicates the experiment and affects the accuracy of the NN. Instead of choosing a specific model, the BNN considers a probability distribution function over model space. It uses the data and Bayes' theorem to convert an initial prior distribution for the models to a posterior distribution. This posterior is then used as a weighting function for the predictions made by the network to new inputs using the different models that are under the prior. The SVM provides an alternative approach to generalisation from that applied in the BNN. It uses capacity control through the maximisation of the margin between the support vectors to improve generalisation performance, however a validation set is still required to select necessary parameters. The main difference between the three techniques and the naive Bayesian classifier is that the latter cannot model correlation between inputs, since it assumes that the observable variables are independent given the class variable. Kernel density estimation that models the class-conditional densities for the NBC is not constrained to any particular functional form.

The four techniques are found to be highly-accurate in classifying FISH signals into their colour, as well as into real and artifact. Frameworks to

combine these classifications into one system are under investigation. The slight difference in the performance of the BNN and the NN can be attributed to the finite (although large) number of samples leading to different maximum likelihood and maximum posterior solutions. The inferiority of the NBC compared with the other techniques can be attributed to the assumption of conditional independence and to the additional inherent feature extraction stage performed by the other classifiers. However, this inferiority should be confronted against the simplicity offered by the NBC. Moreover, if the features are known to be independent from each other, the simple NBC should be our first choice.

Finally, this research can be extended by evaluating other classifiers for the same data and by comparing the classification techniques on other artificial and real-world databases.

Acknowledgements: The authors would like to thank Seema Dhanjal for data collection and signal labelling and William F. Clocksin for helpful discussions. This work is supported by EPSRC contract GR/L51072: *Automatic Analysis of FISH Images*.

References

- [1] Tanke HJ, Florijn RJ, Wiegant J, Raap AK, and Vrolijk J. CCD microscopy and image analysis of cells and chromosomes stained by fluorescence in situ hybridization. *Histochemical Journal*, 27:4–14, 1995.
- [2] Netten H, van Vliet LJ, Vrolijk H, Sloos WCR, Tanke HJ, and Young IT. Fluorescent dot counting in interphase cell nuclei. *Bioimaging*, 4:93–106, 1996.
- [3] Lerner B, Clocksin WF, Dhanjal S, Hultén MA, and Bishop CM. Automatic signal classification in fluorescence in-situ hybridization images. Technical Report 466, University of Cambridge Computer Laboratory, May 1999.
- [4] Lerner B, Clocksin WF, Dhanjal S, Hultén MA, and Bishop CM. Feature representation for the automatic analysis of fluorescence in-situ hybridization images. Technical Report 464, University of Cambridge Computer Laboratory, May 1999.
- [5] Bishop CM. *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford, 1995.
- [6] Cox RT. Probability, frequency and reasonable expectation. *American*

Journal of Physics, 14:1–13, 1946.

- [7] MacKay DJC. A practical Bayesian framework for back-propagation networks. *Neural Computation*, 4:448–472, 1992.
- [8] Gull SF. Bayesian inductive inference and maximum entropy. In Erickson GJ and Smith CR, editors, *Maximum-Entropy and Bayesian Methods in Science and Engineering*, volume 1: Foundations, pages 53–74. Kluwer, Dordrecht, The Netherlands, 1988.
- [9] Hinton GE and van Camp D. Keeping neural networks simple by minimizing the description length of the weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory*, pages 5–13, 1993.
- [10] Barber D and Bishop CM. Ensemble learning in Bayesian neural networks. In Bishop CM, editor, *Neural Networks and Machine Learning*, volume 168 of *Series F: Computer and Systems Sciences*, pages 215–237, Cambridge, U.K., 1998. NATO Advanced Study Institute, Springer-Verlag.
- [11] Lawrence ND and Azzouzi M. A variational Bayesian committee of neural networks. Technical report, University of Cambridge Computer Laboratory,

<http://www.cl.cam.ac.uk/users/nd121/nnmixture.ps.gz>,
1999.

- [12] Neal RM. *Bayesian Learning for Neural Networks*. Springer, 1996. Lecture Notes in Statistics 118.
- [13] Burges CJC. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.
- [14] Haykin S. *Neural Networks: A Comprehensive Foundation*. Macmillan, New York, 2nd edition, 1999.
- [15] Vapnik VN. *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, New York, 1982.
- [16] Vapnik VN and Chervonenkis AY. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- [17] Heckerman D. A tutorial on learning with Bayesian networks. Technical report MSR-TR-95-06, Microsoft Research, March 1995.
- [18] Lerner B. A Bayesian methodology and probability density estimation for fluorescence in-situ hybridization signal classification. Technical Report 474, University of Cambridge Computer Laboratory, October 1999.

- [19] Lerner B, Dhanjal S, and Hultén MA. GELFISH- graphical environment for labelling FISH images. Technical Report 465, University of Cambridge Computer Laboratory, May 1999.
- [20] Neal RM. Software for flexible Bayesian modeling and Markov chain sampling, 1999. available from <http://www.cs.utoronto.ca/~radford/fbm.software.html>.
- [21] Joachims T. Making large-scale svm learning practical. In Schölkopf B, Burges CJC, and Smola AJ, editors, *Advances in Kernel Methods: Support Vector Learning*, pages 169–184, Cambridge, MA, 1998. MIT.
- [22] John GH and Langley P. Estimating continuous distributions in Bayesian classifiers. In P. Besnard and S. Hanks, editors, *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 338–345, San Francisco, CA: Morgan Kaufmann Publishers, 1995.

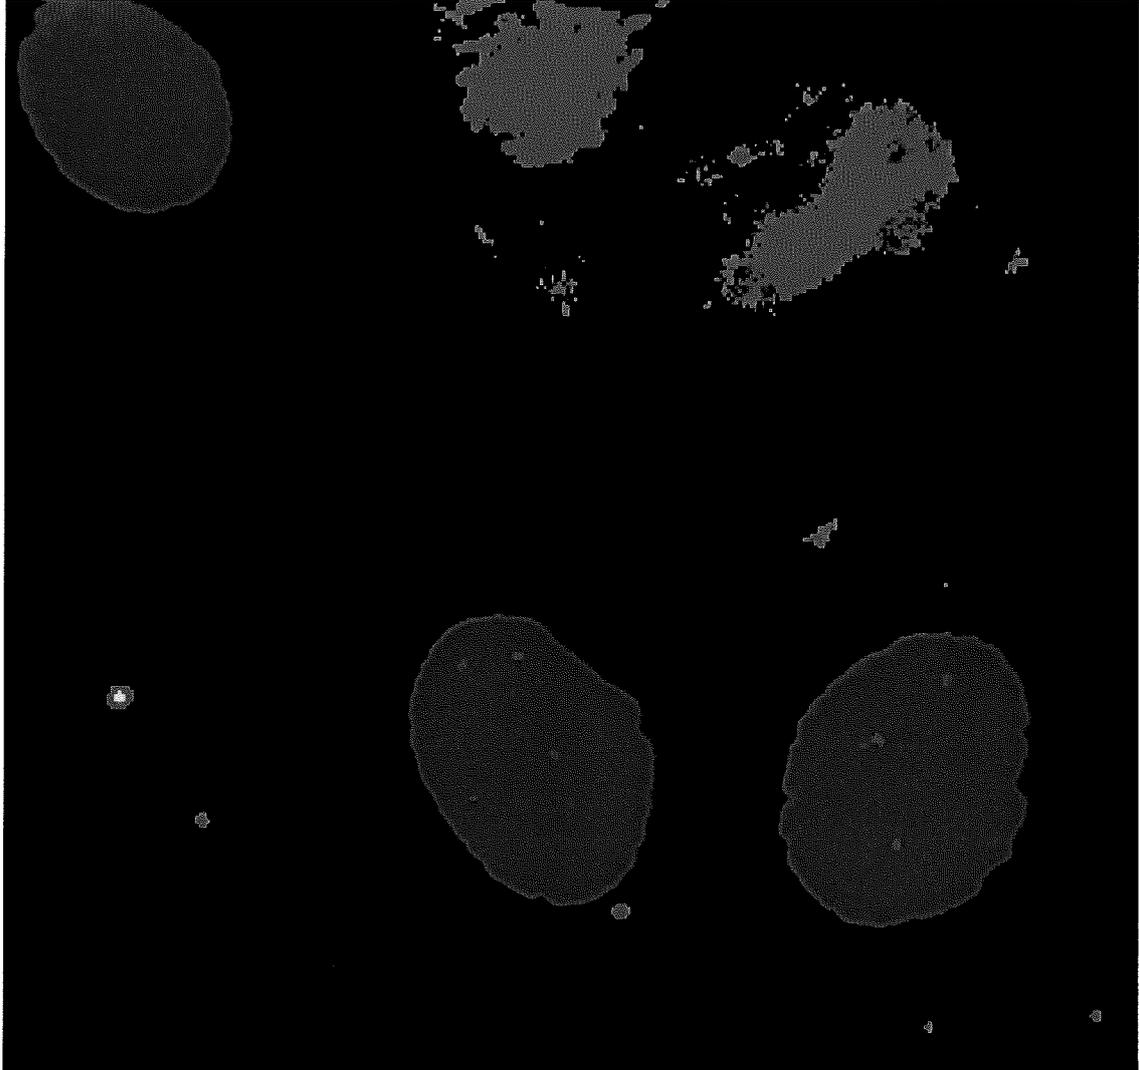


Figure 1: An example of a FISH image used for dot counting

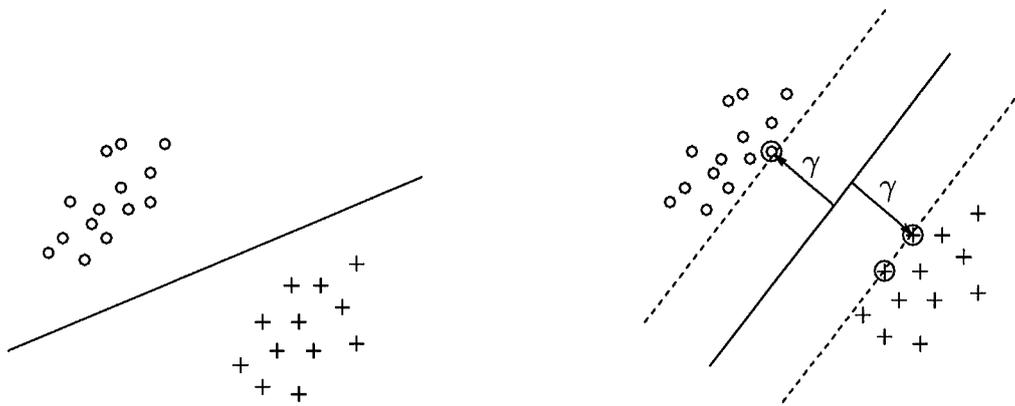


Figure 2: Two different separating hyper-planes which classify all training examples correctly. The diagram on the right shows the maximum margin. The margin size is denoted by γ . Circled data points lie on the margin. They are used to define the classifier and are known as support vectors.

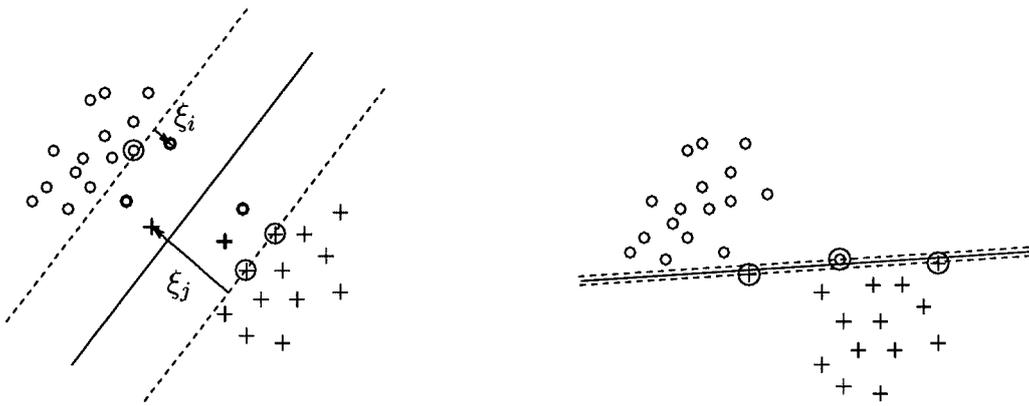


Figure 3: Depending on the penalty assigned to errors, C , different solutions will be found. The left hand figure is a low error penalty, the right hand figure is a high error penalty. Bold data points are within the margin. The classifier pays a penalty, $C\xi_i$, for these data points

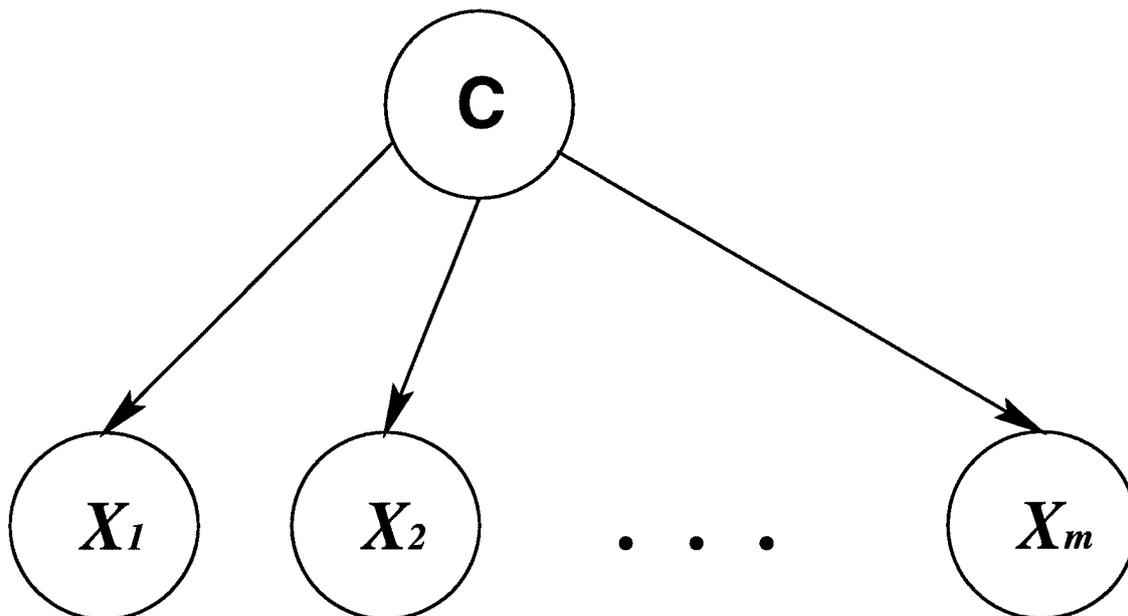


Figure 4: The naive Bayesian classifier depicted as a Bayesian network in which the observable variables (X_1, X_2, \dots, X_m) are conditionally independent given the class variable (C).

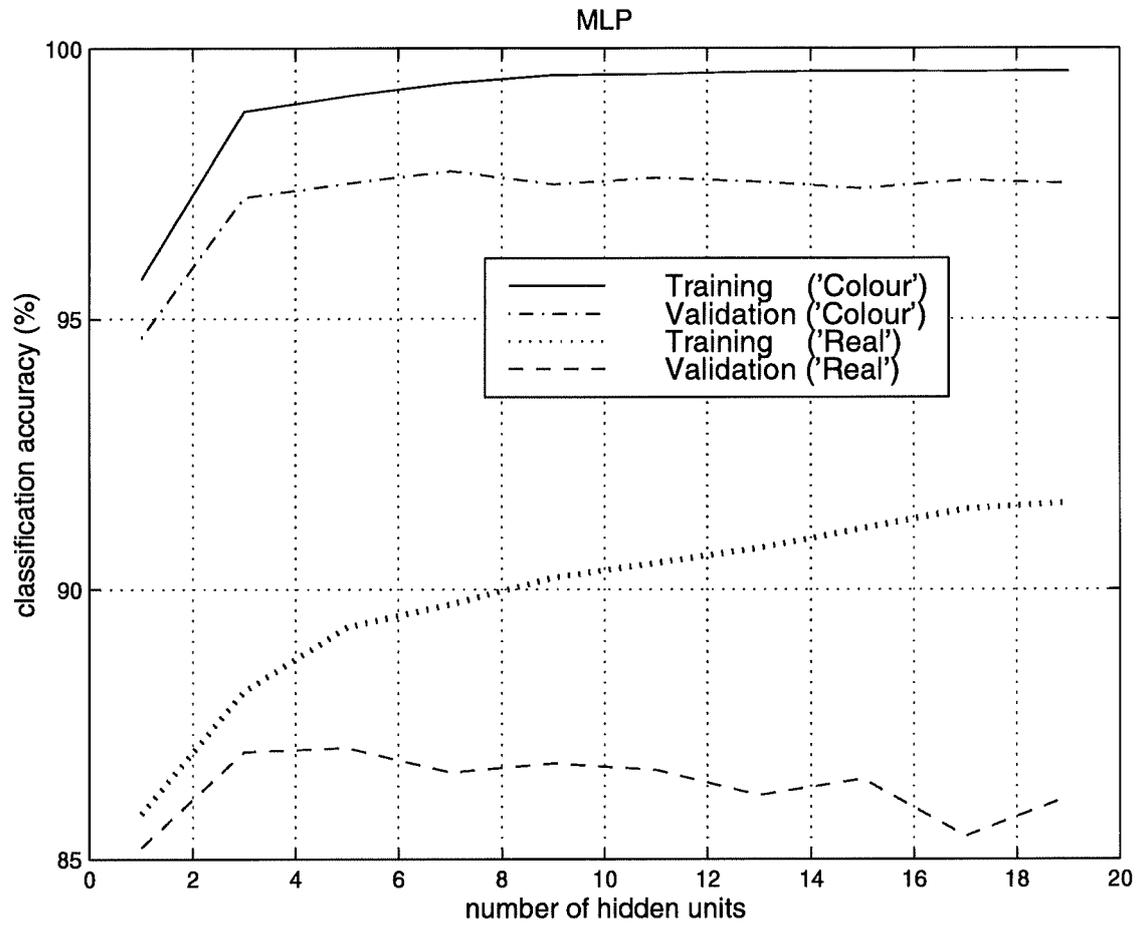


Figure 5: The classification accuracy measured on the training and validation sets of the two MLP NNs – colour NN and real NN for increasing numbers of hidden units. Model selection is based on the highest accuracy on the validation set.

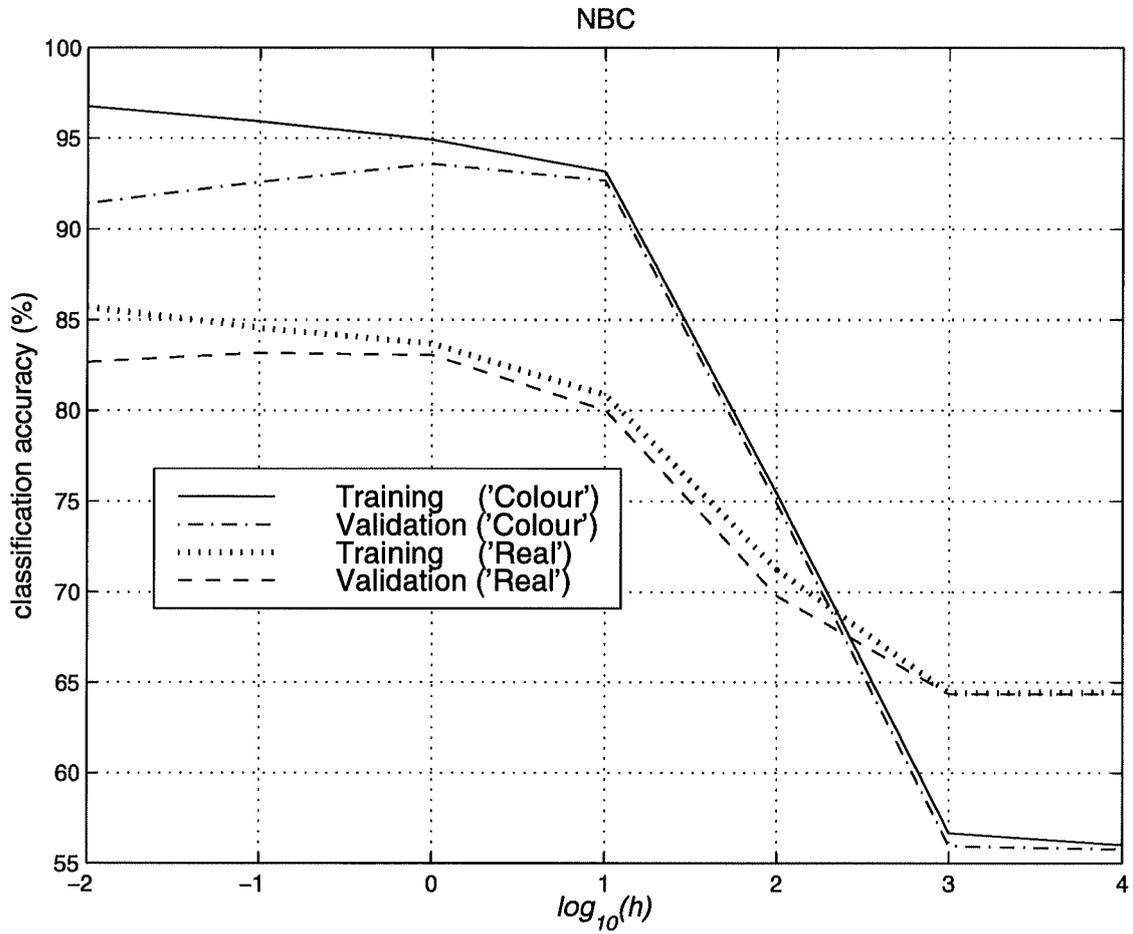


Figure 6: The classification accuracy measured on the training and validation sets of the two NBCs – colour NBC and real NBC for increasing kernel width parameters (h) (on a log scale). Model selection is based on the highest accuracy on the validation set.

Number	Feature	Number	Feature
1	Area	7	Average Hue
2	Eccentricity	8	Hue Texture
3	Total Channel Intensity	9	Delta Hue
4	Average Channel Intensity	10	Eig. 1
5	Texture	11	Eig. 2
6	Maximum Hue	12	Average Grey Intensity

Table 1: The set of features studied in the work. Texture indicates standard deviation of the channel intensity (5) or hue (8). Eig. 1, 2 are abbreviations for the two coordinates of the eigenvector corresponding to the largest eigenvalue of the red and green intensity components of the signal.

model	Real/Artifact (%)	Colour (%)
Neural Network (NN)	86.4	98.1
Bayesian Neural Network (BNN)	88.2	98.8
Support Vector Machine (SVM)	87.2	98.4
Naive Bayesian Classifier (NBC)	83.0	94.0

Table 2: Classification accuracy of the four techniques on the FISH test set.