CAMBRIDGE UNIVERSITY

ENGINEERING DEPARTMENT

SEQUENTIAL MONTE CARLO METHODS FOR OPTIMISATION OF NEURAL NETWORK MODELS

JFG de Freitas, M Niranjan,
AH Gee and A Doucet

CUED/F-INFENG/TR 328

November 18, 1998

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: `jfgf@eng.cam.ac.uk`
URL: `http://www-svr.eng.cam.ac.uk/~jfgf`

Abstract

We discuss a novel strategy for training neural networks using sequential Monte Carlo algorithms and propose a new hybrid gradient descent/sampling importance resampling algorithm (HySIR). In terms of both computational time and accuracy, the hybrid SIR is a clear improvement over conventional sequential Monte Carlo techniques. The new algorithm may be viewed as a global optimisation strategy, which allows us to learn the probability distributions of the network weights and outputs in a sequential framework. It is well suited to applications involving on-line, nonlinear and non-Gaussian signal processing. We show how the new algorithm outperforms extended Kalman filter training on several problems. In particular, we address the problem of pricing option contracts, traded in financial markets. In this context, we are able to estimate the one-step-ahead probability density functions of the options prices.

Contents

1 Introduction

Probabilistic modelling coupled with nonlinear function approximators is a powerful tool in solving many real world problems. In the engineering community, there are many examples of problems involving large datasets that have been tackled with nonlinear function approximators such as the multi layer perceptron (MLP) or radial basis functions. Much of the emphasis here is on performance, in the form of accuracy of prediction on unseen data. The tricks required to estimate parameters of very large models and the handling of very large datasets become interesting challenges too. Tools such as cross validation to deal with over-fitting and bootstrap to deal with model uncertainty have proved to be very effective. Neural networks applied to speech recognition (Robinson 1994), and hand written digit recognition (Le Cun, Boser, Denker, Henderson, Howard and Hubbard 1989) are examples of this. Such work has demonstrated that many interesting and hard inference tasks, involving thousands of free parameters, can indeed be solved at a desirable level of performance. On the other hand, in the statistics community, we see that the emphasis is on rigorous mathematical analysis of algorithms, careful model formulation and model criticism. This work has led to a rich paradigm to handle various inference problems in which a good collection of powerful algorithms exist. Sampling techniques applied to neural networks, starting from the work of Radford Neal is a classic example where both the above features are exploited (Neal 1996). He shows that nonlinear function approximators in the form of MLPs can be trained, and their performance evaluated, in a Bayesian framework. This formulation leads to probability distributions that are difficult to handle analytically. Markov Chain Monte Carlo (MCMC) sampling methods are used to make inferences in the Bayesian framework. While sampling methods tend to be much richer in exploring the probability distributions, approximate methods, such as Gaussian approximation, have also attracted interest (Bishop 1995, Mackay 1992).

Many problems, such as time series analysis, are characterised by data that arrives sequentially. In such problems one has the task of performing model estimation, model validation and inference sequentially, *i.e.* at the arrival of each item of data. By the very nature of such real time tasks, or if we suspect the source of the data to be time varying, we may wish to adopt a sequential strategy in which each item of data is used only once and then discarded. In this case, deciding what information we should be propagating on a sample by sample basis becomes an interesting topic. In the classic state space model estimation by a Kalman filter setting, one defines a Gaussian probability density over the parameters to be estimated. The mean and covariance matrix are propagated using recursive update equations each time new data is received. Use of the matrix inversion lemma allows an elegant update of the inverse covariance matrix without actually having to invert a matrix at each sample. This is a widely studied topic, optimal in the case of linear Gaussian models. For nonlinear models, the common trick is Taylor series expansion around the operating point, leading to the extended Kalman filter (EKF). The derivation of the Kalman filter starting from a Bayesian setting is of much interest. This allows for tuning of the noise levels as well as running multiple models in parallel to evaluate model likelihoods. Jazwinski (Jazwinski 1970) and Bar-Shalom (Bar-Shalom and Li 1993) are classic text books dealing with these aspects. Our earlier work reviews these concepts and shows them applied in a neural network context (de Freitas, Niranjan and Gee 1997, de Freitas, Niranjan and Gee 1998a).

Taylor series approximation leading to the EKF makes gross simplification of the probabilistic specification of the model. With nonlinear models the probability distributions of interest tend to be multi modal. Gaussian approximation in such cases could miss the rich structure brought in by the nonlinear model. Sequential Sampling techniques provide a class of algorithms to address this issue. A good application of these ideas is tracking in computer vision (Isard and Blake 1996), where the Kalman filter approximation is shown to fail to capture the multi modalities. Sequential sampling algorithms, under the names of particle filters, sequential sampling-importance resampling (SIR), bootstrap filters and condensation trackers have also been applied to a wide range of problems, including target tracking (Gordon, Salmond and Smith 1993), financial analysis (Müller 1992, Pitt and Shephard 1997), diagnostic measures of fit (Pitt and Shephard 1997), sequential imputation in missing data problems (Kong, Liu and Wong 1994), blind deconvolution (Liu and Chen 1995) and medical prognosis (Berzuini, Best, Gilks and Larizza 1997).

In this paper we focus on neural networks, constraining ourselves to sequential training and predictions. We use sampling techniques in a state space setting to show how a multi layer perceptron may be trained in environments where data arrives one at a time. We consider sequential importance sampling and resampling algorithms and illustrate their performance on some simple problems. Further, we introduce a hybrid algorithm in which each sampling step is supplemented by a gradient type update step. This could either be a straightforward gradient descent on all the samples propagated, or even an EKF type update of the samples. It should be noted that we do not deal with the problem of choosing the network architecture.

We target this paper primarily at the neural networks community. We believe many signal processing applications addressed by this community will fall into this category. Part of the paper is written in a tutorial form to introduce the concepts involved here. We also believe the real life example included here, and the experience gained in applying sequential sampling techniques with a highly nonlinear function approximator will be of benefit to the statistics community. In particular, we improve the existing sequential sampling methods by incorporating gradient information. This results in a hybrid optimisation scheme (HySIR), whereby each sampling trajectory of the sampling algorithm is updated with and extended Kalman filter. The HySIR may, therefore, be interpreted as an efficient dynamic mixture of extended Kalman filters, whose accuracy improves as the number of filters increases. We show in this paper that the HySIR algorithm outperforms conventional sequential sampling methods in terms of both computational speed and accuracy.

In Section 2, we formulate the problem of training neural networks in terms of a state space representation. A theoretical Bayesian solution is subsequently proposed in Section 3. Owing to the practical difficulties that arise when computing the Bayesian solution, three approximate methods (numerical integration, Gaussian approximation and Monte Carlo simulation) are discussed in Section 4. In Sections 5 and 6, we derive a generic sequential importance sampling methodology and point out some of the limitations of conventional sampling methods. In section 7, we introduce the HySIR algorithm. Finally, in Section 8, we test the algorithms on several problems including the pricing of options on the FTSE-100 index.

2 State Space Neural Network Modelling

As in our previous work (de Freitas et al. 1997), we start from a state space representation to model the neural network's evolution in time. A transition equation describes the evolution of the network weights, while a measurements equation describes the nonlinear relation between the inputs and outputs of a particular physical process, as follows:

$$\mathbf{w}_{k+1} = \mathbf{w}_k + \mathbf{d}_k \tag{1}$$

$$\mathbf{y}_k = \mathbf{g}(\mathbf{w}_k, \mathbf{x}_k) + \mathbf{v}_k \tag{2}$$

where ($\mathbf{y}_k \in \Re^o$) denotes the output measurements, ($\mathbf{x}_k \in \Re^d$) the input measurements and ($\mathbf{w}_k \in \Re^m$) the neural network weights. The measurements nonlinear mapping $\mathbf{g}(.)$ is approximated by a multi-layer perceptron (MLP). It is widely known that this neural model exhibits the capacity to approximate any continuous function, to an arbitrary precision, as long as it is not restricted in size (Cybenko 1989). Yet, our work may be easily extended to encompass recurrent networks, radial basis networks and many other approximation techniques.

The measurements are assumed to be corrupted by noise $\mathbf{v}_k$. In the sequential Monte Carlo framework, the probability distribution of the noise is specified by the user. In our examples we shall choose a zero mean Gaussian distribution with covariance R. The measurement noise is assumed to be uncorrelated with the network weights and initial conditions.

We model the evolution of the network weights by assuming that they depend on the previous value $\mathbf{w}_k$ and a stochastic component $\mathbf{d}_k$. The process noise $\mathbf{d}_k$ may represent our uncertainty in how the parameters evolve, modelling errors or unknown inputs such as manoeuvres in tracking applications (Bar-Shalom and Li 1993). We assume the process noise to be a zero mean Gaussian process with covariance Q, however other distributions can also be adopted. This choice of distributions for the network weights requires further research. The process noise is also assumed to be uncorrelated with the network weights. In Section 8, we propose that this random walk model can be improved by replacing it with the following gradient descent model:

$$\mathbf{w}_{k+1} = \mathbf{w}_k + \alpha(\mathbf{y}_k - \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k}))\frac{\partial \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k})}{\partial \mathbf{w}_\mathrm{k}} + \mathbf{d}_\mathrm{k} \tag{3}$$

where α is a learning rate parameter and $\hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k})$ is the model prediction at time k. The gradient descent model allows for directed steps when searching for optimal locations in parameter space.

The noise terms are assumed to be uncorrelated with the network weights and the initial conditions. The evolution of the states (network weights) corresponds to a Markov process with initial probability $\mathrm{p}(\mathbf{w}_0)$ and transition probability $\mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1})$. The observations are assumed to be conditionally independent given the states. These are standard assumptions in a large class of tracking and time series problems (Harvey 1970).

Typically, the task is to estimate the network weights $\hat{\mathbf{w}}_k$ and the noise parameters R and Q given the measurements $Y_k = \{\mathbf{y}_1,\ \mathbf{y}_2,\ \cdots,\ \mathbf{y}_k\}$. To simplify the exposition in this paper, we do not treat the problem of estimating the noise covariances and the initial probabilities. To understand how these variables may be estimated via hierarchical Bayesian models or EM learning, the reader is referred to our previous work (de Freitas et al. 1997, de Freitas et al. 1998a). We shall, however, analyse the roles played by R and Q in sequential Monte Carlo simulation.

3 The Bayesian Solution

The posterior density $\mathrm{p}(W_k|Y_k)$, where $Y_k = \{\mathbf{y}_1,\ \mathbf{y}_2,\ \cdots,\ \mathbf{y}_k\}$ and $W_k = \{\mathbf{w}_1,\ \mathbf{w}_2,\ \cdots,\ \mathbf{w}_k\}$, constitutes the complete solution to the sequential estimation problem. In many applications, such as tracking, it is of interest to estimate one of its marginals, namely the filtering density $\mathrm{p}(\mathbf{w}_k|Y_k)$. By computing the filtering density recursively, we do not need to keep track of the complete history of the weights. Thus, from a storage point of view, the filtering density turns out to be more parsimonious than the full posterior density function. If we know the filtering density of the network weights, we can easily derive various estimates of the network weights, including centroids, modes, medians and confidence intervals. In the following section, we show how the filtering density may be approximated using sequential importance sampling techniques.

The filtering density is estimated recursively in two stages: prediction and update, as illustrated in Figure 1. In the prediction step, the filtering density $\mathrm{p}(\mathbf{w}_{k-1}|Y_{k-1})$ is propagated into the future via the transition density $\mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1})$ as follows:

$$\mathrm{p}(\mathbf{w}_k|Y_{k-1}) = \int \mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1})\mathrm{p}(\mathbf{w}_{k-1}|Y_{k-1})\mathrm{d}\mathbf{w}_{k-1} \tag{4}$$

The transition density is defined in terms of the probabilistic model governing the states'

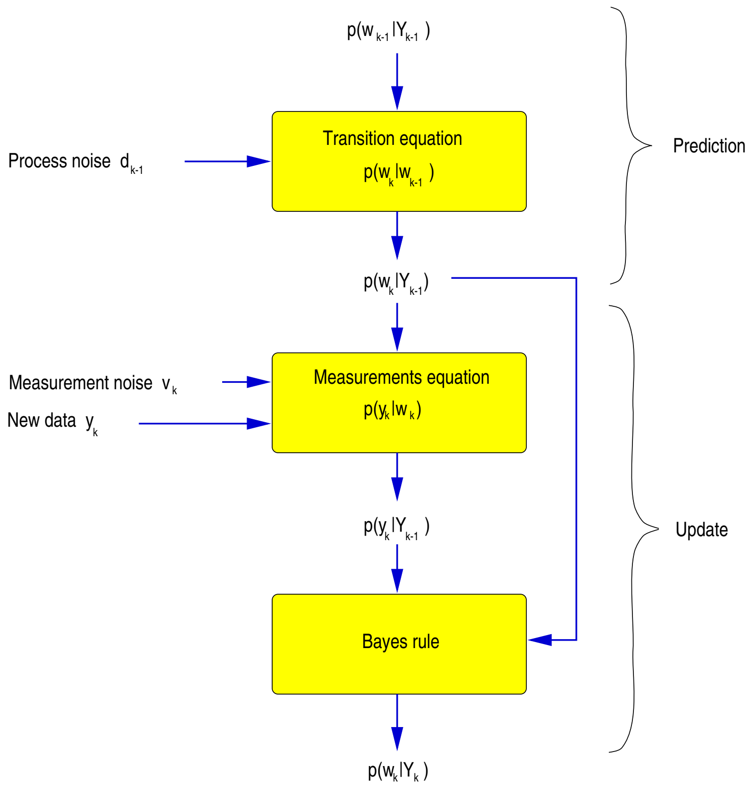

Figure 1: Prediction and update stages in the recursive computation of the filtering density.

evolution (equation (1)) and the process noise statistics. That is:

$$\begin{aligned}
\mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1}) &= \int \mathrm{p}(\mathbf{w}_k|\mathbf{d}_{k-1}, \mathbf{w}_{k-1})\mathrm{p}(\mathbf{d}_{k-1}|\mathbf{w}_{k-1})\mathrm{d}\mathbf{d}_{k-1} \\
&= \int \delta(\mathbf{w}_k - \mathbf{d}_{k-1} - \mathbf{w}_{k-1})\mathrm{p}(\mathbf{d}_{k-1})\mathrm{d}\mathbf{d}_{k-1}
\end{aligned}$$

where the Dirac delta function $\delta(.)$ indicates that $\mathbf{w}_k$ can be computed via a purely deterministic relation when $\mathbf{w}_{k-1}$ and $\mathbf{d}_{k-1}$ are known. Note that $\mathrm{p}(\mathbf{d}_{k-1}|\mathbf{w}_{k-1}) = \mathrm{p}(\mathbf{d}_{k-1})$ because the process and measurement noise terms have been assumed to be independent of past and present values of the states.

The update stage involves the application of Bayes' rule when new data is observed:

$$\mathrm{p}(\mathbf{w}_k|Y_k) = \frac{\mathrm{p}(\mathbf{y}_k|\mathbf{w}_k)\mathrm{p}(\mathbf{w}_k|Y_{k-1})}{\mathrm{p}(\mathbf{y}_k|Y_{k-1})} \tag{5}$$

The likelihood density function is defined in terms of the measurements model (equation (2)) as follows:

$$\mathrm{p}(\mathbf{y}_k|\mathbf{w}_k) = \int \delta(\mathbf{y}_k - g(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k}) - \mathbf{v}_\mathrm{k})\mathrm{p}(\mathbf{v}_\mathrm{k})\mathrm{d}\mathbf{v}_\mathrm{k}$$

The normalising denominator of equation (5) is often referred to as the evidence density function. It plays a key role in optimisation schemes that exploit Gaussian approximation (de Freitas et al. 1997, Jazwinski 1969, Mackay 1992, Sibisi 1989). It is given by:

$$\mathrm{p}(\mathbf{y}_k|Y_{k-1}) = \int \mathrm{p}(\mathbf{y}_k|\mathbf{w}_k)\mathrm{p}(\mathbf{w}_k|Y_{k-1})\mathrm{d}\mathbf{w}_k$$

In this formulation, the question of how to choose the prior $p(\mathbf{w}_0)$ arises immediately. This is a central issue in Bayesian inference. The problem is, however, much more general. It is related to the ill-posed nature of the inference and learning task. A finite set of data admits many possible explanations, unless some form of *a priori* knowledge is used to constrain the solution space. Assuming that the data conforms to a specific degree of smoothness is certainly the most fundamental and widely applied type of *a priori* knowledge. This assumption is particularly suited to applications involving noisy data. Nonetheless, the Bayesian approach admits more general forms of *a priori* knowledge within a probabilistic formulation. It would be interesting to devise methods for mapping experts' knowledge, available in many formats, into the Bayesian probabilistic representation.

4 Practical Solutions

The prediction and update strategy given by equations (4) and (5) provides an optimal solution to the inference problem, but, unfortunately, it involves multi-dimensional integrations. These integrations are the source of most of the practical difficulties inherent to the Bayesian methodology. In most applications, analytical solutions are not possible. This is why we need to resort to alternative approaches, such as direct numerical integration, Gaussian approximation and Monte Carlo simulation methods.

4.1 Direct Numerical Integration

The direct numerical integration method relies on approximating the distribution of interest by a discrete distribution on a finite grid of points. The location of the grid is a non-trivial design issue. Once the distribution is computed at the grid points, an interpolation procedure is used to approximate it in the remainder of the space. Kitagawa (Kitagawa 1987) used this method to replace the filtering integrals by finite sums over a large set of equally spaced grid points. He chose a piece-wise linear interpolation strategy. Kramer and Sorenson (Kramer and Sorenson 1988) adhered to the same methodology, but opted for a constant interpolating function. Pole and West (Pole and West 1990) have attempted to reduce the problem of choosing the grid's location by implementing a dynamic grid allocation method.

When the grid points are spaced "closely enough" and encompass the region of high probability, the method works well. However, the method is very difficult to implement in high-dimensional multivariate problems such as neural network modelling. Here, computing at every point in a dense multi-dimensional grid becomes prohibitively expensive (Gelman, Carlin, Stern and Rubin 1995, Gilks, Richardson and Spiegelhalter 1996).

4.2 Gaussian Approximation

Until recently, the popular approach to sequential estimation has been Gaussian approximation (Bar-Shalom and Li 1993). In the linear Gaussian estimation problem, the Kalman filter provides an optimal recursive algorithm for propagating and updating the mean and covariance of the hidden states. In nonlinear scenarios, the EKF is a computationally efficient natural extension of the Kalman filter. It is based on a Taylor expansion of the dynamics and measurements nonlinear equations about the last predicted state. Typically, first order expansions, which lead to a Gaussian approximation of the posterior density function, are employed. The mean and covariance are propagated and updated by a simple set of equations, similar to the Kalman filter equations. As the number of terms in the Taylor expansion increases, the complexity of the algorithm also increases due to the computation of derivatives of increasing order. For example, a linear expansion requires the computation of the Jacobian, while a quadratic expansion involves computing the Hessian matrix.

Gaussian approximation results in a simple and elegant framework that is amenable to the design of inference and learning algorithms. In the past, various authors have studied the problem of approximating the distribution of neural network weights with a Gaussian function. One of the earliest implementations of EKF trained MLPs is due to Singhal and Wu (Singhal and Wu 1988). In their method, the network weights are grouped into a single vector $\mathbf{w}$ that is updated in accordance with the EKF equations. The entries of the Jacobian matrix (*i.e.* the derivative of the outputs with respect to the weights) are calculated by back-propagating the output observations through the network.

The algorithm proposed by Singhal and Wu requires a considerable computational effort. The complexity is of the order om^2 multiplications per estimation step, where m represents the number of weights and o the number of network outputs. Shah, Palmieri and Datum (Shah, Palmieri and Datum 1992) and Puskorius and Feldkamp (Puskorius and Feldkamp 1991, Puskorius, Feldkamp and Davis 1996) have proposed strategies for decoupling the global EKF estimation algorithm into local EKF estimation sub-problems. For example, they suggest that the weights of each neuron could be updated independently. The assumption in the local updating strategies is that the weights are decoupled and, consequently, the Kalman filter weights covariance is a block-diagonal matrix.

The EKF is an improvement over conventional neural network estimation techniques, such as on-line back-propagation, in that it makes use of second order statistics (covariances). These statistics are essential for placing error bars on the predictions and for combining separate networks into committees of networks. The back-propagation algorithm is, in fact, a degenerate form of the EKF algorithm (de Freitas et al. 1997, Ruck, Rogers, Kabrisky, Maybeck and Oxley 1992). Previously, we have shown that Gaussian approximation, within a hierarchical Bayesian framework, leads to interesting algorithms that exhibit adaptive memory and adaptive regularisation coefficients. These algorithms enabled us to learn the dynamics and measurements noise statistics in slowly changing non-stationary environments (de Freitas et al. 1997, de Freitas, Niranjan and Gee 1998c). In stationary environments, where the data is available in batches, we have also shown in (de Freitas et al. 1998a, de Freitas, Niranjan and Gee 1998b) that it is possible to learn these statistics by means of the Rauch-Tung-Striebel smoother (Rauch, Tung and Striebel 1965) and the EM algorithm (Dempster, Laird and Rubin 1977). There have been several attempts to implement the EM algorithm on-line (Collings, Krishnamurthy and Moore 1994, Elliott 1994, Krishnamurthy and Moore 1993). Unfortunately, these still rely on various heuristics and tend to be computationally demanding.

A natural progression on sequential Gaussian approximation is to employ a mixture of Gaussian densities (Kadirkamanathan and Kadirkamanathan 1995, Li and Bar-Shalom 1994, Sorenson and Alspach 1971). These mixtures can either be static or dynamic. In static mixtures, the Gaussian models assumed to be valid throughout the entire process are a subset of several hypothesised models. That is, we start with a few models and compute which of them describe the sequential process most accurately. The remaining models are, subsequently, discarded. In dynamic model selection, one particular model out of a set of r operating models is selected during each estimation step. Dynamic mixtures of Gaussian models are far more general than static mixtures of models. However, in stationary environments, static mixtures are obviously more adequate. Dynamic mixtures are particularly suited to the problem of noise estimation in rapidly changing environments, such as tracking manoeuvring targets. There each model corresponds to a different hypothesis of the value of the noise covariances (Bar-Shalom and Li 1993).

In summary, Gaussian approximation, because of its simplicity and computational efficiency, constitutes a good way of handling many problems where the density of interest has a significant and predominant mode. Many problems, however, do not fall into this category. Mixtures of Gaussians provide a better solution when there are a few dominant modes. However, they introduce extra computational requirements and complications, such as estimating the number of mixture components.

4.3 Monte Carlo Simulation

Many signal processing problems, such as equalisation of communication signals, financial time series, medical prognosis, target tracking and geophysical data analysis, involve elements of non-Gaussianity, nonlinearity and non-stationarity. Consequently, it is not possible to derive exact closed form estimators based upon the standard criteria of maximum likelihood, maximum *a posteriori* or minimum mean-squared error. Analytical approximations to the true distribution of the data do not take into account all the salient statistical features of the processes under consideration. Monte Carlo simulation methods, which provide a complete description of the probability distribution of the data, improve the accuracy of the analysis.

The basic idea in Monte Carlo simulation is that a set of weighted samples, drawn from the posterior density function of the neural network weights, is used to map the integrations, involved in the inference process, to discrete sums. More precisely, we make use of the following Monte Carlo approximation:

$$\hat{\mathrm{p}}(W_k|Y_k) = \frac{1}{N}\sum_{i=1}^{N} \delta(W_k - W_k^{(i)})$$

where $W_k^{(i)}$ represents the samples used to describe the posterior density and, as before, $\delta(.)$ denotes the Dirac delta function. We carry the suffix k, denoting time, to emphasise the fact that the approximation is performed at at the arrival of each data point. Consequently, any expectations of the form:

$$\mathbf{E}[f_k(W_k)] = \int f_k(W_k)\mathrm{p}(W_k|Y_k)\mathrm{d}W_k$$

may be approximated by the following estimate:

$$\mathbf{E}[f_k(W_k)] \approx \frac{1}{N}\sum_{i=1}^{N} f_k(W_k^{(i)})$$

where the samples $W_k^{(i)}$ are drawn from the posterior density function. Monte Carlo sampling techniques are an improvement over direct numerical approximation in that they automatically select samples in regions of high probability.

In recent years, many researchers in the the statistical and signal processing communities have, almost simultaneously, proposed several variations of sequential Monte Carlo algorithms. These algorithms have been applied to a wide range of problems, including target tracking (Gordon et al. 1993, Isard and Blake 1996), financial analysis (Müller 1992, Pitt and Shephard 1997), diagnostic measures of fit (Gelfand, Dey and Chang 1992, Pitt and Shephard 1997), sequential imputation in missing data problems (Kong et al. 1994), blind deconvolution (Liu and Chen 1995) and medical prognosis (Berzuini et al. 1997). Most of the current research has focused on statistical refinements of this class of algorithm (Carpenter, Clifford and Fearnhead 1997, Liu and Chen 1998, Pitt and Shephard 1997). The basic sequential Monte Carlo methods had already been introduced in the automatic control field in the late sixties. For instance, Handschin and Mayne (Handschin and Mayne 1969) tackled the problem of nonlinear filtering with a sequential importance sampling approach. They combined analytical and numerical techniques, using the Control Variate Method (Hammersley and Handscomb 1968), to reduce the variance of the estimates. In the seventies, various researchers continued working on these ideas (Akashi and Kumamoto 1977, Handschin 1970, Zaritskii, Svetnik and Shimelevich 1975). The paper of Zaritskii, Svetnik and Shimelevich (Zaritskii et al. 1975) is particularly interesting. They treated the problems of continuous and discrete time filtering and introduced several novel ideas. With the exception of Doucet (Doucet 1998), most authors have overlooked the Monte Carlo sampling ideas proposed in the seventies.

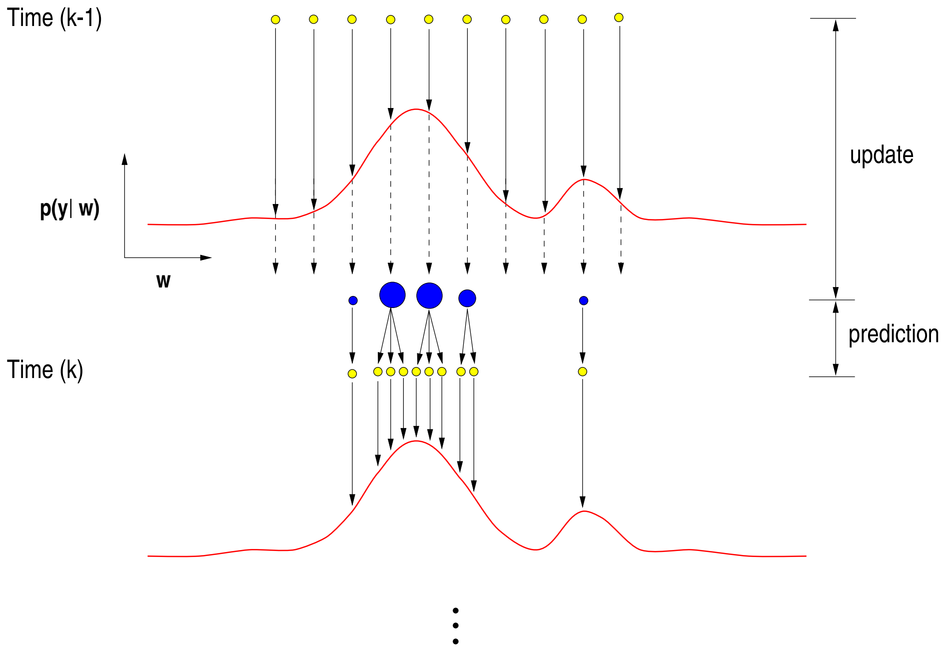

Figure 2: Update and prediction stages in the sequential sampling process. In the update stage, the likelihood of each sample is evaluated. The samples are then propagated with their respective likelihood. The size of the black circles indicates the likelihood of a particular sample. The samples with higher likelihood are allowed to have more "children". In the prediction stage, a process noise term is added to the samples so as to increase the variety of the sample set. The end result is that the surviving samples provide a better weighted description of the likelihood function.

Figure 2 illustrates the operation of a generic sequential Monte Carlo method. It embraces the standard assumption that we can sample from the prior $p(\mathbf{w}_k|Y_{k-1})$ and evaluate the likelihood $p(\mathbf{y}_k|\mathbf{w}_k^{(i)})$ of each sample. Only the fittest samples, that is the ones with the highest likelihood, survive after the update stage. These then proceed to be multiplied, according to their likelihood, in the prediction stage. As mentioned earlier, the update and prediction stages are governed by equations (5) and (4) respectively. It is instructive to approach the problem from an optimisation perspective. Figures 3 and 4 show the windowed global descent in the error function that is typically observed. The diagrams shed light on the roles played by the noise covariances R and Q. Q dictates by how much the cloud of samples is expanded in the prediction stage. By increasing Q, the density of the cloud of samples is reduced. Consequently, the algorithm will take longer to converge. A very small Q, on the other hand, will not allow the algorithm to explore new regions of the parameter space. Ideally, one needs to implement an algorithm that adapts Q automatically. This is explored in Section 7. R controls the resolution of the update stage, as shown in Figure 5. A small value of R will cause the likelihood to be too narrow. Consequently, only a few trajectories will be able to propagate to the next time step. If R is too small, the optimisation scheme might fail to detect some of the important modes of the likelihood function. By increasing R, we broaden the likelihood function, thereby increasing the number of surviving trajectories. If we choose R to be too large, all the trajectories become equally likely and the algorithm might not converge. As shown in Figures 3 and 4, R gives rise to a threshold $T(R)$ in the error function, which determines the width of the updated clouds of samples.

If we can sample from the prior and evaluate the likelihood up to proportionality, three

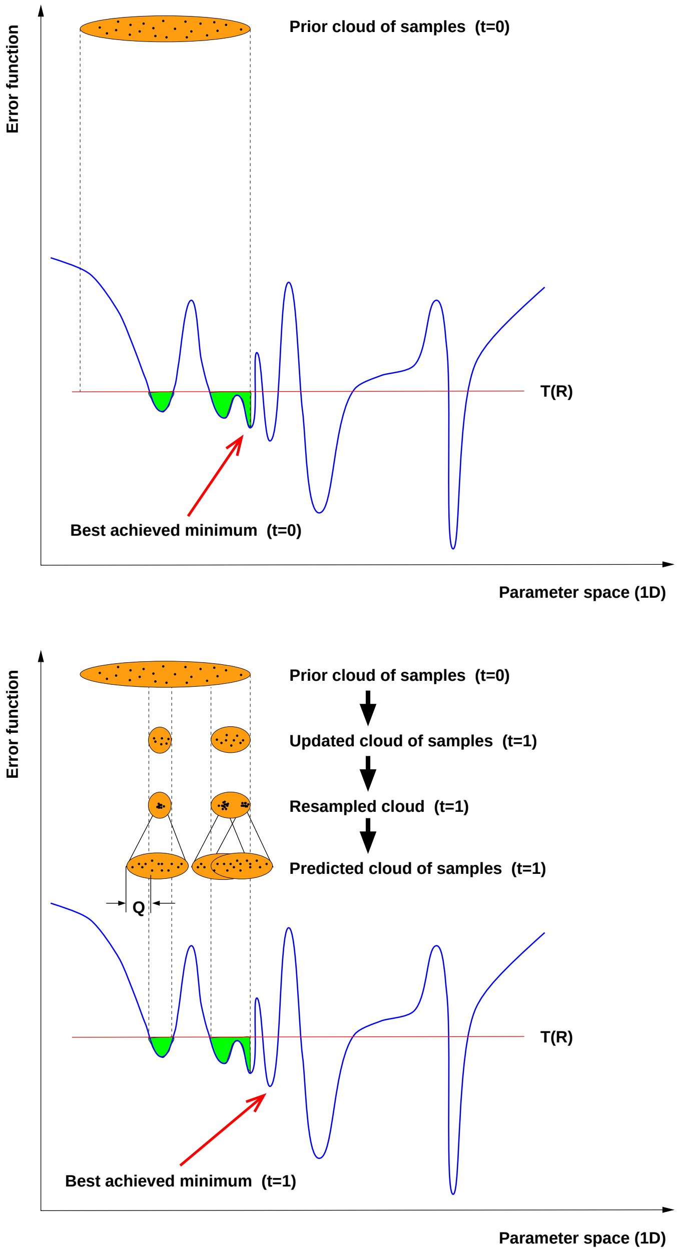

Figure 3: First and second steps of a one-dimensional sequential Monte Carlo optimisation problem. The size of the initial cloud of samples determines the search region in parameter space. The goal is to find the best possible minimum of the error function. It is clear that as the number of samples increases, our chances of reaching lower minima increase. In the second step, the updated clouds of samples are generated. The width of these clouds is obtained from the intersection between the width of the prior cloud of samples, the error function and a threshold determined by the measurements noise covariance R. The updated clouds of samples are denser than the prior cloud of samples. Next, the samples are grouped in regions of higher likelihood in a resampling step. Finally, the clouds of samples are expanded by a factor determined by the process noise covariance Q. This expansion allows the clouds to reach regions of the parameter function where the error function is lower.

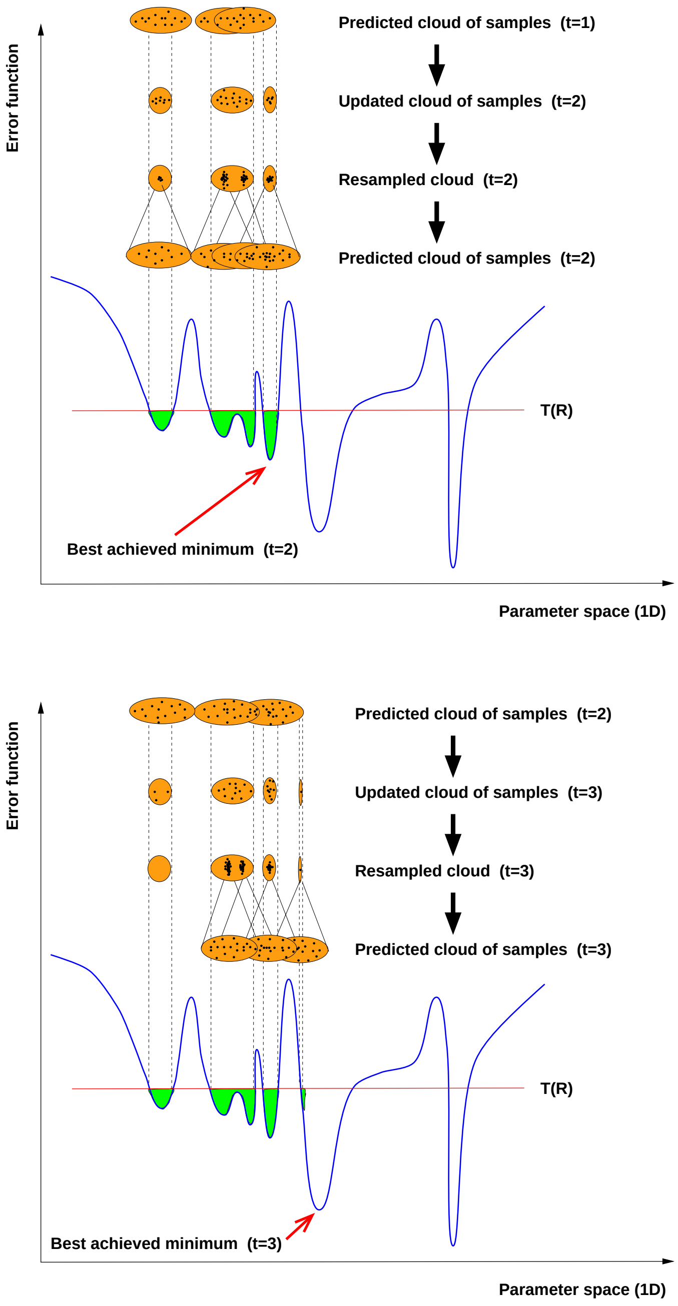

Figure 4: Third and fourth step of a one-dimensional sequential Monte Carlo optimisation problem. To reach the global minimum on the right, the number of samples has to be increased.

different sampling strategies can be adopted to sample from the posterior, namely acceptance-rejection sampling, Markov Chain Monte Carlo (MCMC) and sampling-importance resampling (SIR).

After sampling from the prior, rejection sampling dictates that one should accept the samples with probability $p(\mathbf{y}_k|\mathbf{w}_k^{(i)})/p(\mathbf{y}_k|\mathbf{w}_{k-Max}^{(i)})$, where $\mathbf{w}_{k-Max}^{(i)} = \arg\max_{\mathbf{w}_k} p(\mathbf{y}_k|\mathbf{w}_k)$. Unfortunately, the rejection sampler requires a random number of iterations at each time step.

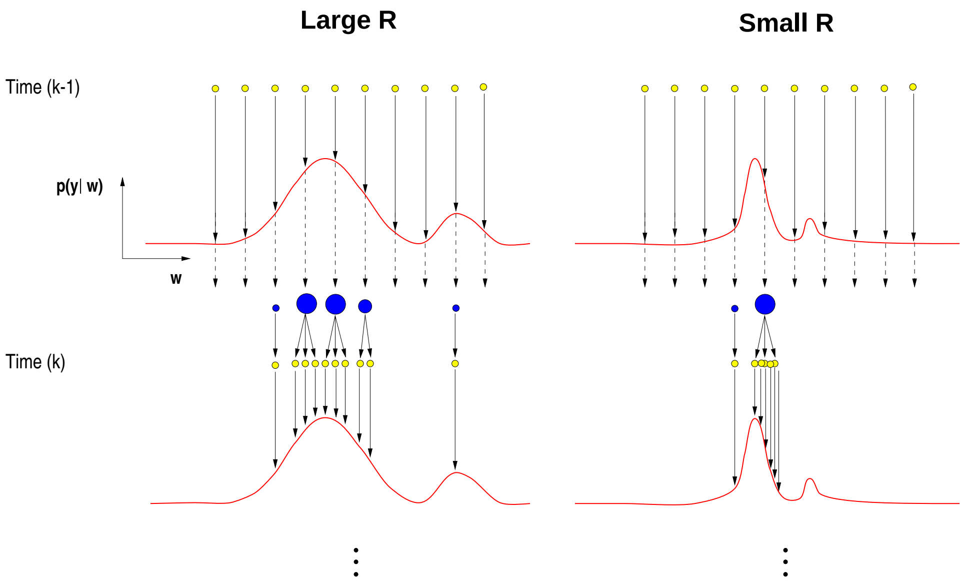

Figure 5: Role played by the measurements variance R. Small values of R result in a narrow likelihood, with only a few trajectories being able to propagate to the next time step. If R is too small, we might miss some of the important modes. On the other hand, if R is too large, we are not giving preference to any of the trajectories and the algorithm might not converge.

This proves to be computationally expensive in high-dimensional spaces (Doucet 1998, Müller 1991, Pitt and Shephard 1997).

The fundamental idea behind MCMC methods is to construct a Markov chain, whose asymptotic distribution tends to the required posterior density (Gilks et al. 1996). Generally, it can take a large number of iterations before this happens. Moreover, it is difficult to even assess when the chain has converged. For most real-time sequential processing applications, MCMC methods can be too computationally demanding (Berzuini et al. 1997, Gordon and Whitby 1995, Müller 1992).

In our work, we favour the SIR sampling option. This approach is the core of many successful sequential Bayesian tools, such as the well known bootstrap or particle filters (Avitzour 1995, Beadle and Djuric 1997, Berzuini et al. 1997, Doucet 1998, Gordon et al. 1993, Kitagawa 1996, Liu and Chen 1998, Pitt and Shephard 1997). In the following sections, we derive this algorithm from an importance sampling perspective.

5 Importance Sampling

As mentioned in Section 4.3, we can approximate the posterior density function with a function on a finite discrete support. Consequently, if follows from the strong law of large numbers that as the number or samples N increases, expectations can be mapped into sums. Unfortunately, it is often impossible to sample directly from the posterior density function. However, we can circumvent this difficulty by sampling from a known, easy-to-sample, proposal density function $\pi(W_k|Y_k)$ and making use of the following substitution:

$$\begin{aligned}\mathbf{E}[f_k(W_k)] &= \int f_k(W_k)\frac{\mathrm{p}(W_k|Y_k)}{\pi(W_k|Y_k)}\pi(W_k|Y_k)\mathrm{d}W_k \\ &= \int f_k(W_k)\frac{\mathrm{p}(Y_k|W_k)\mathrm{p}(W_k)}{\mathrm{p}(Y_k)\pi(W_k|Y_k)}\pi(W_k|Y_k)\mathrm{d}W_k\end{aligned}$$

$$= \int f_k(W_k)\frac{q_k(W_k)}{\mathrm{p}(Y_k)}\pi(W_k|Y_k)\mathrm{d}W_k$$

where the variables $q_k(W_k)$ are known as the unnormalised importance ratios:

$$q_k = \frac{\mathrm{p}(Y_k|W_k)\mathrm{p}(W_k)}{\pi(W_k|Y_k)} \tag{6}$$

We can get rid of the unknown normalising density $\mathrm{p}(Y_k)$ as follows:

$$\begin{aligned}
\mathbf{E}[f_k(W_k)] &= \frac{1}{\mathrm{p}(Y_k)}\int f_k(W_k)q_k(W_k)\pi(W_k|Y_k)\mathrm{d}W_k \\
&= \frac{\int f_k(W_k)q_k(W_k)\pi(W_k|Y_k)\mathrm{d}W_k}{\int \mathrm{p}(Y_k|W_k)\mathrm{p}(W_k)\frac{\pi(W_k|Y_k)}{\pi(W_k|Y_k)}\mathrm{d}W_k} \\
&= \frac{\int f_k(W_k)q_k(W_k)\pi(W_k|Y_k)\mathrm{d}W_k}{\int q_k(W_k)\pi(W_k|Y_k)\mathrm{d}W_k} \\
&= \frac{\mathbf{E}_\pi[q_k(W_k)f_k(W_k)]}{\mathbf{E}_\pi[q_k(W_k)]}
\end{aligned}$$

Hence, by drawing samples from the proposal function $\pi(.)$, we can approximate the expectations of interest by the following estimate:

$$\begin{aligned}
\mathbf{E}[f_k(W_k)] &\approx \frac{1/N\sum_{i=1}^N f_k(W_k^{(i)})q_k(W_k^{(i)})}{1/N\sum_{i=1}^N q_k(W_k^{(i)})} \\
&= \sum_{i=1}^N f_k(W_k^{(i)})\tilde{q}_k(W_k^{(i)})
\end{aligned} \tag{7}$$

where the normalised importance ratios $\tilde{q}_k^{(i)}$ are given by:

$$\tilde{q}_k^{(i)} = \frac{q_k^{(i)}}{\sum_{j=1}^N q_k^{(j)}}$$

The estimate of equation (6) is valid if the $W_k^{(i)}$ correspond to a set of i.i.d. samples drawn from the proposal density function, the support of the proposal function includes the support of the posterior function and the expectations of $f(.)$, q and $qf^2(.)$ exist and are finite (Geweke 1989). Under these conditions and as N tends to infinity, the posterior density function can be approximated arbitrarily well by the estimate:

$$\hat{\mathrm{p}}(W_k|Y_k) = \sum_{i=1}^N \tilde{q}_k^{(i)}\delta(W_k - W_k^{(i)})$$

6 Sequential Importance Sampling

In order to compute a sequential estimate of the posterior density function at time k without modifying the previously simulated states W_{k-1}, we may adopt the following proposal density:

$$\begin{aligned}
\pi(W_k|Y_k) &= \pi(W_0|Y_0)\prod_{j=1}^k \pi(\mathbf{w}_j|W_{j-1}, Y_j) \\
&= \pi(W_{k-1}|Y_{k-1})\pi(\mathbf{w}_k|W_{k-1}, Y_k)
\end{aligned} \tag{8}$$

At this stage, we need to recall that we assumed that the states correspond to a Markov process and that the observations are conditionally independent given the states. That is:

$$\mathrm{p}(W_k) = \mathrm{p}(\mathbf{w}_0) \prod_{j=1}^{k} \mathrm{p}(\mathbf{w}_j|\mathbf{w}_{j-1}) \tag{9}$$

and

$$\mathrm{p}(Y_k|W_k) = \prod_{j=1}^{k} \mathrm{p}(\mathbf{y}_j|\mathbf{w}_j) \tag{10}$$

By substituting equations (7), (9) and (10) into equation (8), we can derive a recursive estimate for the importance weights as follows:

$$\begin{aligned}
q_k &= \frac{\mathrm{p}(Y_k|W_k)\mathrm{p}(W_k)}{\pi(W_{k-1}|Y_{k-1})\pi(\mathbf{w}_k|W_{k-1},Y_k)} \\
&= q_{k-1}\frac{\mathrm{p}(Y_k|W_k)\mathrm{p}(W_k)}{\mathrm{p}(Y_{k-1}|W_{k-1})\mathrm{p}(W_{k-1})}\frac{1}{\pi(\mathbf{w}_k|W_{k-1},Y_k)} \\
&= q_{k-1}\frac{\mathrm{p}(\mathbf{y}_k|\mathbf{w}_k)\mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1})}{\pi(\mathbf{w}_k|W_{k-1},Y_k)} \qquad (11)
\end{aligned}$$

Equation (11) gives a mechanism to sequentially update the importance ratios. Since we can sample from the proposal function and evaluate the likelihood and transition probabilities (equations (1) and (2)), all we need to do is generate a prior set of samples and iteratively compute the importance ratios. This procedure allows us to obtain the following type of estimates:

$$\mathbf{E}[f_k(W_k)] \approx \sum_{i=1}^{N} \tilde{q}_k^{(i)} f_k(W_k^{(i)}) \tag{12}$$

In Section 6.4 we give a more detailed description of a generic sequential importance sampling algorithm. We need to delay this presentation because the SIS algorithm discussed so far has a serious limitation that requires our immediate attention: the variance of the ratio $(\mathrm{p}(W_k|Y_k)/\pi(W_k|Y_k))$ increases over time when the observations are regarded as random (Kong et al. 1994). To understand why the variance increase poses a problem, suppose that we want to sample from the posterior. In that case, we want the proposal density to be very close to the posterior density. When this happens, we obtain the following results for the mean and variance:

$$\mathbf{E}_\pi \left[\frac{\mathrm{p}(W_k|Y_k)}{\pi(W_k|Y_k)} \right] = 1$$

and

$$var_\pi \left[\frac{\mathrm{p}(W_k|Y_k)}{\pi(W_k|Y_k)} \right] = \mathbf{E}_\pi \left[\left(\frac{\mathrm{p}(W_k|Y_k)}{\pi(W_k|Y_k)} - 1\right)^2 \right] = 0$$

In other words, we expect the variance to be close to zero in order to obtain reasonable estimates. Therefore, a variance increase has a harmful effect on the accuracy of the simulations. In practice, the degeneracy caused by the variance increase can be observed by monitoring the importance ratios. Typically, what we observe is that, after a few iterations, one of the normalised importance ratios tends to 1, while the remaining ratios tend to zero.

6.1 Resampling

To avoid the degeneracy of the SIS simulation method, a resampling stage may be used to eliminate samples with low importance ratios and multiply samples with high importance ratios. It is possible to see an analogy to the steps in genetic algorithms. Many of the ideas on resampling have stemmed from the work of Efron (Efron 1982), Rubin (Rubin 1988) and Smith and Gelfand (Smith and Gelfand 1992). Various authors have described efficient algorithms for accomplishing this task in $\mathcal{O}(N)$ operations (Doucet 1998, Pitt and Shephard 1997).

Resampling involves mapping the measure $\{\mathbf{w}_k^{(i)}, \tilde{q}_k^{(i)}\}$ into an equally weighted random measure $\{\mathbf{w}_k^{(j)}, N^{-1}\}$. This is accomplished by sampling uniformly from the discrete set $\{\mathbf{w}_k^{(i)}\}_{i=1}^N$ with probabilities $\{\tilde{q}_k^{(i)}\}_{i=1}^N$. A mathematical proof of this can be found on pages 111–112 of (Gordon 1994). Figure 6 shows a way of sampling from this discrete set. After constructing the cumulative distribution of the discrete set, a uniformly drawn sampling index i is projected onto the the distribution range and then onto the distribution domain. The intersection with the domain constitutes the new sample index j. That is, the vector $\mathbf{w}_k^{(j)}$ is accepted as the new sample. Clearly, the vectors with the larger sampling ratios will end up with more copies after the resampling process. In the subsequent prediction stage, random disturbances are added to the samples, thereby adding variety to the dominant samples, as was shown in Figure 2.

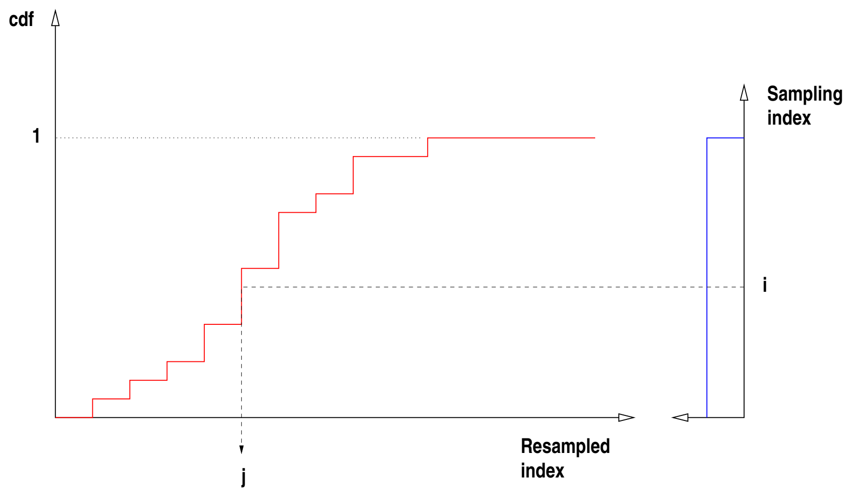

Figure 6: Resampling process, whereby a random measure $\{\mathbf{w}_k^{(i)}, \tilde{q}_k^{(i)}\}$ is mapped into an equally weighted random measure $\{\mathbf{w}_k^{(j)}, N^{-1}\}$. The index i is drawn from a uniform distribution.

Liu and Chen (Liu and Chen 1995, Liu and Chen 1998) have argued that when all the importance ratios are nearly equal, resampling only reduces the number of distinctive streams and introduces extra variation in the simulations. Using a result from (Kong et al. 1994), they suggest that resampling should only be performed when the effective sample size N_{eff} is below a fixed heuristic threshold. The effective sample size is defined as:

$$N_{eff} = \frac{1}{\sum_{i=1}^N \left(\tilde{q}_k^{(i)} \right)^2}$$

We have so far explained how to compute the importance ratios sequentially and how to improve the sample set by resampling. After discussing the choices of likelihood and proposal density functions, we shall be able to present a complete and detailed algorithm description.

6.2 Likelihood Function

The likelihood density function is defined in terms of the measurements equation. In our work, we adopted the following Gaussian form:

$$\mathrm{p}(\mathbf{y}_k|\mathbf{w}_k) \propto \exp -\frac{1}{2}\left((\mathbf{y}_k - \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k},\mathbf{x}_\mathrm{k}))^\mathrm{T}\mathrm{R}^{-1}(\mathbf{y}_\mathrm{k} - \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k},\mathbf{x}_\mathrm{k}))\right)$$

We also studied the effect of adding a weight decay term to the likelihood function to penalise for large network weights. This did not seem to improve the results in the experiments that we performed. In addition, it introduced an extra tuning parameter. If it is suspected that the data set may contain several outliers, likelihood functions with heavy tails should be employed.

6.3 Choosing the Proposal Function

The choice of proposal function is one of the most critical design issues in importance sampling algorithms. Doucet (Doucet 1997) has shown that the proposal function:

$$\pi(\mathbf{w}_k|W_{k-1},Y_k) = \mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1},\mathbf{y}_k) \tag{13}$$

minimises the variance of the importance ratios conditional on W_{k-1} and Y_k. This choice of proposal function has also been advocated by other researchers (Kong et al. 1994, Liu and Chen 1995, Zaritskii et al. 1975). Nonetheless, the density:

$$\pi(\mathbf{w}_k|W_{k-1},Y_k) = \mathrm{p}(\mathbf{w}_k|\mathbf{w}_{k-1}) \tag{14}$$

is the most popular choice of proposal function (Avitzour 1995, Beadle and Djuric 1997, Gordon et al. 1993, Isard and Blake 1996, Kitagawa 1996). Although it results in higher Monte Carlo variation than the optimal proposal (equation (13)), as a result of it not incorporating the most recent observations, it is usually easier to implement (Berzuini et al. 1997, Doucet 1998, Liu and Chen 1998).

In Section 7, we describe an algorithm that allows us to update each sampling trajectory, with an extended Kalman filter, before the resampling stage. This Kalman update constitutes, effectively, a strategy for sampling from the prior of equation (13).

6.4 Algorithm Pseudo-code

The essential structure of a sequential Monte Carlo algorithm to train neural networks with two layers, using the proposal of function of equation (14), involves the following steps:

1. `INITIALISE NETWORK WEIGHTS` $(k = 0)$:

   `For` $i = 1, \cdots, N$

   - `Draw the weights` $\mathbf{w}_0^{(i)}$ `from the first layer prior` $\mathrm{p}_1(\mathbf{w}_0)$ `and the second layer prior` $\mathrm{p}_2(\mathbf{w}_0)$.
   - `Evaluate the importance ratios:`

     $$q_0^{(i)} = \mathrm{p}(\mathbf{y}_0|\mathbf{w}_0^{(i)})$$

   - `Normalise the importance ratios:`

     $$\tilde{q}_0^{(i)} = \frac{q_0^{(i)}}{\sum_{j=1}^{N} q_0^{(j)}}$$

2. For $k = 1, \cdots, L$

   (a) SAMPLING STAGE:

   For $i = 1, \cdots, N$

   - Predict via the dynamics equation:

   $$\hat{\mathbf{w}}_{k+1}^{(i)} = \mathbf{w}_k^{(i)} + \mathbf{d}_k^{(i)}$$

     where $\mathbf{d}_k^{(i)}$ is a sample from $\mathrm{p}(\mathbf{d}_k)$ ($\mathcal{N}(0, Q)$ in our case).

   - Evaluate the importance ratios:

   $$q_{k+1}^{(i)} = q_k^{(i)} \mathrm{p}(\mathbf{y}_{k+1} | \hat{\mathbf{w}}_{k+1}^{(i)})$$

   - Normalise the importance ratios:

   $$\tilde{q}_{k+1}^{(i)} = \frac{q_{k+1}^{(i)}}{\sum_{j=1}^{N} q_{k+1}^{(j)}}$$

   (b) RESAMPLING STAGE:

   For $i = 1, \cdots, N$

   If $N_{eff} \geq$ Threshold:

   - $\mathbf{w}_{k+1}^{(i)} = \hat{\mathbf{w}}_{k+1}^{(i)}$

   Else

   - Resample new index j from the discrete set $\{\hat{\mathbf{w}}_{k+1}^{(i)}, \tilde{q}_{k+1}^{(i)}\}$
   - $\mathbf{w}_{k+1}^{(i)} = \hat{\mathbf{w}}_{k+1}^{(j)}$
   - $q_{k+1}^{(i)} = \frac{1}{N}$

An order N algorithm to perform the resampling stage is described in pseudo-code in a paper by Carpenter, Clifford and Fearnhead (Carpenter et al. 1997). The generic sequential sampling algorithm is rather straightforward to implement using a Gaussian likelihood function[1]. We make use of a different prior for each layer because the functional form of the neurons varies with layer (Müller and Rios Insua 1998). In addition, we can select the prior hyper-parameters for each layer in accordance with the magnitude of the input and output signals. Statistical estimates such as mean values and confidence intervals can be computed using equation (12).

6.5 Problems with Sequential SIR Simulation

The discreteness of the resampling stage (see Figure 6) implies that any particular sample with a high importance ratio will be duplicated many times. As a result, the cloud of samples may eventually collapse to a single sample. This degeneracy, will limit the ability of the algorithm to search for lower minima in other regions of the error surface. In addition, the number of samples used to describe the posterior density function will become too small and inadequate. Gordon (Gordon 1994) discusses various strategies to overcome this problem. He suggests that the sample set can be boosted by ensuring that the number of elements in the discrete distribution exceeds the posterior samples by a factor of at least ten. In addition, he proposes smoothing the discrete distribution with a kernel interpolator. In another paper (Gordon et al. 1993), Gordon, Salmond and Smith propose a roughening procedure, whereby

[1] We have made available the software for the implementation of the SIR algorithm at the following website: `http://svr-www.eng.cam.ac.uk/~jfgf/software.html`.

an independent jitter, with zero mean and standard deviation σ, is added to each sample drawn in the resampling stage. The standard deviation of the jitter is given by:

$$\sigma = KIN^{-1/m} \tag{15}$$

where I denotes the length of the interval between the maximum and minimum samples of each specific component, K is a tuning parameter and, as before, m represents the number of weights. Large values of K blur the distribution, while very small values produce tight clusters of points around the original samples. We have found, in our simulations, that this technique works very well and that tuning K requires little effort.

Another serious problem arises when the likelihood function is much narrower than the prior. In this scenario, very few samples are accepted in the resampling stage. The problem is exacerbated when the regions of high likelihood are located in the tails of the prior density. This difficulty may be surmounted by the brute force method of increasing the sample size. Alternatively, one should select the prior variances and likelihood variance with great care, after an exploratory data analysis stage.

In our neural networks framework, we adopt different informative priors for each layer, whose variances depend on the magnitude of the input and output signals. Moreover, the covariance matrix of the network weights for each sampling trajectory is propagated and updated with an extended Kalman filter. This leads to an improved annealed sampling algorithm, whereby the variance of each sampling trajectory decreases with time. That is, we start searching over a large region of the error surface and as time progresses, we concentrate on the regions of lower error. This technique is described in the following section.

7 HySIR

The Monte Carlo conception of optimisation relies solely on probing the error surface at several points as shown in Figures 3 and 4. It fails to make use of all the information implicit in the error surface. For instance, the method could be enhanced by evaluating the gradient and other higher derivatives of the error surface. This is evident in Figure 7. By allowing the samples to follow the gradient after the first prediction stage, the algorithm can reach a better minimum.

In order to improve existing sequential Monte Carlo simulation algorithms, we propose a new hybrid gradient descent/SIR method. The main feature of the algorithm is that the samples are updated by a gradient descent step in the prediction stage, as follows:

$$\begin{aligned}\mathbf{w}_{k+1} &= \mathbf{w}_k - \frac{\alpha}{2}\frac{\partial}{\partial \mathbf{w}_k}(\mathbf{y}_k - \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k}))^2 \\ &= \mathbf{w}_k + \alpha(\mathbf{y}_k - \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k}))\frac{\partial \hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k})}{\partial \mathbf{w}_\mathrm{k}} \end{aligned} \tag{16}$$

The term $\partial\hat{\mathbf{g}}(\mathbf{w}_\mathrm{k}, \mathbf{x}_\mathrm{k})/\partial\mathbf{w}_\mathrm{k}$ is the Jacobian. It can be easily computed by back-propagating derivatives through the network as explained in Appendix B of (de Freitas et al. 1997). The gradient descent learning rate α is a tuning parameter that controls how fast we should descend. Too large a parameter would cause the trajectory in the error surface to over-fluctuate, thereby never converging. A very small value, on the other hand, would not contribute towards speeding the SIR algorithm's convergence.

The plain gradient descent algorithm can get trapped at shallow local minima. One way of overcoming this problem is to define the error surface in terms of a Hamiltonian that accounts for the approximation errors and the momentum of each trajectory. This is the basis of the Hybrid Monte Carlo algorithm (Brass, Pendleton, Chen and Robson 1993, Duane, Kennedy, Pendleton and Roweth 1987). In this algorithm, each trajectory is updated by approximating the Hamiltonian differential equations by a leapfrog discretisation scheme. The

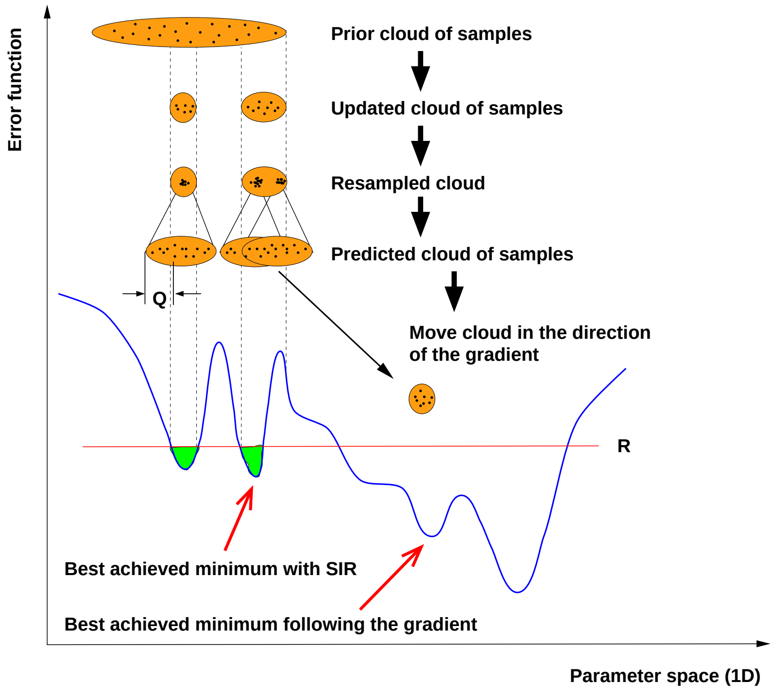

Figure 7: HySIR algorithm. By following the gradient after the prediction stage, it is possible to reach better minima.

discretisation may, however, introduce biases. In our work, we favour a stochastic gradient descent approach based on the extended Kalman filter. This technique avoids shallow local minima by adaptively adding noise to the network weights. This is shown in Figure 8.

The extended Kalman filter equations are given by:

$$\begin{aligned}
\mathbf{w}_{k+1|k} &= \mathbf{w}_k \\
P_{k+1|k} &= P_k^T + Q^* I_{mm} \\
K_{k+1} &= P_{k+1|k} G_{k+1} [R^* I_{oo} + G_{k+1}^T P_{k+1|k} G_{k+1}]^{-1} \\
\mathbf{w}_{k+1} &= \mathbf{w}_{k+1|K} + K_{k+1}(\mathbf{y}_{k+1} - \mathbf{g}(\mathbf{x}_{\mathrm{k}}, \mathbf{w}_{\mathrm{k+1|k}})) \\
P_{k+1} &= P_{k+1|k} - K_{k+1} G_{k+1}^T P_{k+1|k}
\end{aligned}$$

where K_{k+1} is known as the Kalman gain matrix, I_{mm} denotes the identity matrix of size $m \times m$, and R^* and Q^* are two tuning parameters, whose roles are explained in detail in (de Freitas et al. 1997). Here, it suffices to say that they control the rate of convergence of the EKF algorithm and the generalisation performance of the neural network. In the general multiple input, multiple output (MIMO) case, $\mathbf{g} \in \Re^{\mathrm{o}}$ is a vector function and G represents the Jacobian matrix:

$$G = \frac{\partial \mathbf{g}}{\partial \mathbf{w}} \mid_{(\mathbf{w}=\hat{\mathbf{w}})} = \begin{bmatrix} \frac{\partial \mathbf{g}_1}{\partial \mathbf{w}_1} & \frac{\partial \mathbf{g}_2}{\partial \mathbf{w}_1} & \cdots & \frac{\partial \mathbf{g}_o}{\partial \mathbf{w}_1} \\ \frac{\partial \mathbf{g}_1}{\partial \mathbf{w}_2} & & & \\ \vdots & & & \vdots \\ \frac{\partial \mathbf{g}_1}{\partial \mathbf{w}_m} & \cdots & & \frac{\partial \mathbf{g}_o}{\partial \mathbf{w}_m} \end{bmatrix}^T$$

Since the EKF is a suboptimal estimator based on linearisation of a nonlinear mapping, strictly speaking, P_k is an approximation to the covariance matrix. In mathematical terms:

$$P_k \approx \mathbf{E}[(\mathbf{w}_k - \hat{\mathbf{w}}_k)^T (\mathbf{w}_k - \hat{\mathbf{w}}_k) | Y_k]$$

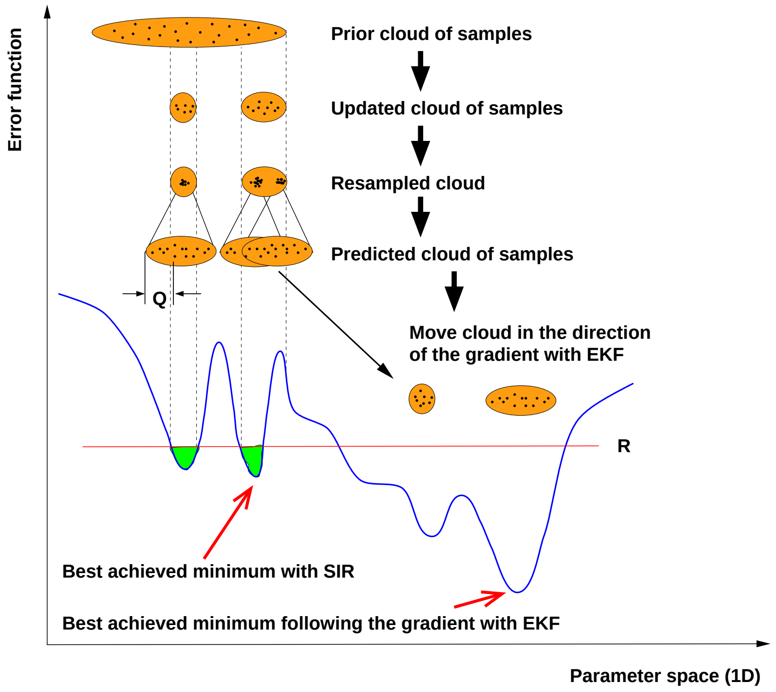

Figure 8: HySIR with EKF updating. By adding noise to the network weights and adapting the weights covariance, it is possible to reach better minima with the EKF than with plain gradient descent techniques.

The Kalman filter step, before the update stage, allows us to incorporate the latest measurements into the prior. As a result, we can sample from the prior $p(\mathbf{w}_k|\mathbf{w}_{k-1}, \mathbf{y}_k)$ before the update stage. Another advantage of this method is that the covariance of the weights changes over time. Because the extended Kalman filter is a minimum variance estimator, the covariance of the weights decreases with time. Consequently, as the tracking improves, the variation of the network weights is reduced. This annealing process improves the efficiency of the search for global minima in parameter space and reduces the variance of the estimates. It should be pointed out that the weights need to be resampled in conjunction with their respective covariances. The pseudo-code for the HySIR algorithm with EKF updating is as follows[2]:

1. `INITIALISE NETWORK WEIGHTS` $(k = 0)$:
2. `For` $k = 1, \cdots, L$
   (a) `SAMPLING STAGE:`
   `For` $i = 1, \cdots, N$
   - `Predict via the dynamics equation:`
   $$\hat{\mathbf{w}}_{k+1}^{(i)} = \mathbf{w}_k^{(i)} + \mathbf{d}_k^{(i)}$$
   `where` $\mathbf{d}_k^{(i)}$ `is a sample from` $p(\mathbf{d}_k)$ $(\mathcal{N}(0, Q)$ `in our case).`
   - `Update samples with the EKF equations.`

[2]We have made available the software for the implementation of the HySIR algorithm at the following web-site: `http://svr-www.eng.cam.ac.uk/~jfgf/software.html`.

- `Evaluate the importance ratios:`

$$q_{k+1}^{(i)} = q_k^{(i)} \mathrm{p}(\mathbf{y}_{k+1}|\hat{\mathbf{w}}_{k+1}^{(i)})$$

- `Normalise the importance ratios:`

$$\tilde{q}_{k+1}^{(i)} = \frac{q_{k+1}^{(i)}}{\sum_{j=1}^{N} q_{k+1}^{(j)}}$$

(b) `RESAMPLING STAGE:`

`For` $i = 1, \cdots, N$

`If` $N_{eff} \geq$ `Threshold:`

- $\mathbf{w}_{k+1}^{(i)} = \hat{\mathbf{w}}_{k+1}^{(i)}$
- $P_{k+1}^{(i)} = \hat{P}_{k+1}^{(i)}$

`Else`

- `Resample new index` j `from the discrete set` $\{\hat{\mathbf{w}}_{k+1}^{(i)}, \tilde{q}_{k+1}^{(i)}\}$
- $\mathbf{w}_{k+1}^{(i)} = \hat{\mathbf{w}}_{k+1}^{(j)}$ `and` $P_{k+1}^{(i)} = \hat{P}_{k+1}^{(j)}$
- $q_{k+1}^{(i)} = \frac{1}{N}$

The HySIR with EKF updates may be interpreted as a dynamical mixture of EKFs. That is, the estimates depend on a few modes of the posterior density function. In addition to having to compute the normalised importance ratios ($\mathcal{O}(N(oS_1^2 + d^2 S_1))$) operations, where S_1 is the number of hidden neurons, N the number of samples, d the input dimension and o the number of outputs) and resampling ($\mathcal{O}(N)$ operations), the HySIR has to perform the extended Kalman filter updates ($\mathcal{O}(Nom^2)$ operations, where m is the number of network weights). That is, the computational complexity of the HySIR increases approximately linearly with the number of EKFs in the mixture. For problems with dominant modes, we only need to use a few trajectories and, consequently, the HySIR provides accurate and computationally efficient solutions.

The HySIR approach makes use of the likelihood function twice per time step. Consequently, if we make the noise distributions too narrow, the algorithm will only give a representation of a narrow region of the posterior density function. Typically, the algorithm will converge to a few modes of the posterior density function. This is illustrated in Section 8.2.

8 Simulations

In this section, we discuss two experiments, using synthetic data, and a real application involving the pricing financial options on the FTSE100 index. The first experiment provides a comparison between the various algorithms discussed in this paper. It addresses the trade-off between accuracy and computational efficiency. The second experiment shows how the HySIR algorithm can be used to estimate time-varying latent variables. It also serves the purpose of illustrating the effect of the process noise covariance Q on the simulations.

8.1 Experiment 1: Time-Varying Function Approximation

We generated input-output data using the following time-varying function:

$$y(x_1(k), x_2(k)) = 4\sin(x_1(k) - 2) + 2x_2(k)^2 + 5\cos(0.02k) + 5 + \nu$$

where the inputs $x_1(k)$ and $x_2(k)$ were simulated from a Gaussian distribution with zero mean and unit variance. The noise ν was drawn from a Gaussian distribution with zero mean and variance equal to 0.1. Subsequently, we approximated the data with an MLP with 5 hidden sigmoidal neurons and a linear output neuron. The MLP was trained sequentially using the HySIR, SIR, SIS and EKF algorithms. The difference between the SIR and the SIS algorithms is that SIR resamples every time, while SIS only resamples if the effective sample set is below a threshold. We set our threshold to $N/3$.

We set the number of samples (N) to 100 for the SIS and SIR methods and to 10 for the HySIR method. We assigned the values 100 and 1 to the initial variance of the weights and the diagonal entries of the weights covariance matrix. The diagonal entries of R and Q were given the values 0.5 and 2 respectively. Finally, the extended Kalman filter noise hyper-parameters R^* and Q^* were set to 2 and 0.01.

Table 1 shows the average one-step-ahead prediction errors obtained for 100 runs, of 200 time steps each, on a Silicon Graphics R10000 workstation. On average, there was a 50% occurrence of resampling steps using the SIS algorithm. It is clear from the results that avoiding the resampling stage does not yield a great reduction in computational time. This is because one needs to evaluate neural network mappings for each trajectory in the sampling stage. As a result the sampling stage tends to be much more computationally expensive than the resampling stage. The results also show that the HySIR performs much better than the

	EKF	**SIS**	**SIR**	**HySIR**
RMS Error	6.51	3.87	3.27	1.17
Mega flops	4.8	5.6	5.8	48.6
Time (seconds)	1.2	27.4	28.9	24.9

Table 1: RMS errors and floating point operations for the algorithms used in the function approximation problem. It should be noted that the RMS errors include the convergence period. This explains why they are higher than the variance of the additive noise.

conventional SIR and EKF algorithms. It is, however, computationally more expensive, as explained at the end of Section 7.

It is interesting to note that the HySIR, despite requiring more floating point operations, is faster than the SIR algorithm. This shows that the computational efficiency of sampling algorithms depends on the platform in which they are implemented. The design of efficient software and hardware platforms is an important avenue for further research.

Figure 9 shows two plots of the prediction errors for the EKF and HySIR. The left plot shows the best and worst results obtained with the EKF out of 100 trials; each trial involving a different initial weights vector. We then used the 100 initial weights vectors to initialise the HySIR algorithm with 100 sampling trajectories. As shown in the right plot, the global search nature of the HySIR algorithm allows it to converge much faster than the EKF algorithm. This is a clear example that dynamic mixtures of models, with model selection at each time step, perform better than static mixtures (choosing the best EKF out of the 100 trials).

8.2 Experiment 2: Latent States Tracking

To assess the ability of the hybrid algorithm to estimate time-varying hidden parameters, we generated input-output data from a logistic function followed by a linear scaling and a displacement as shown in Figure 10. We applied two Gaussian input sequences to the model. This simple model is equivalent to an MLP with one hidden neuron and an output linear neuron. We then trained a second model with the same structure using the input-output data generated by the first model.

In the training phase, of 200 time steps, we allowed the model weights to vary with time. During this phase, the HySIR algorithm was used to track the input-output training data and estimate the latent model weights. After the 200-th time step, we fixed the values of the

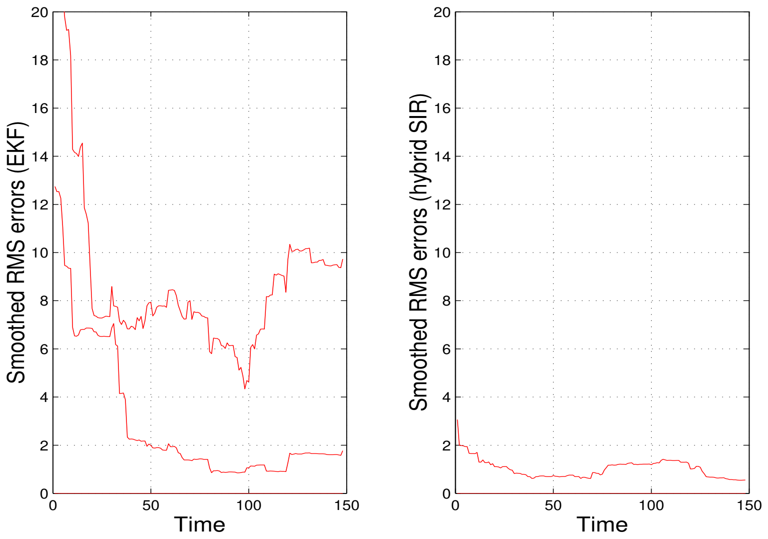

Figure 9: Convergence of the EKF and HySIR algorithms for Experiment 1. The left plot shows the best and worst EKF convergence results out of 100 runs using different initial conditions. The initial conditions for each EKF were used to initialise the EKF filters within the HySIR algorithm. At each time step, the HySIR evaluates a weighted combination of several EKFs. By performing this type of dynamic model selection, it is able to search over large regions of parameter space and, hence, converge much faster than the standard EKF algorithm.

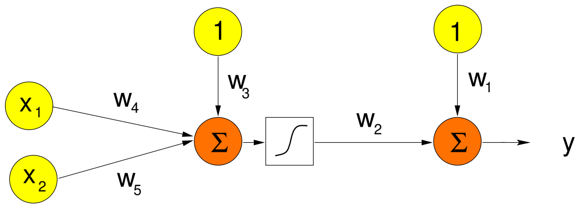

Figure 10: Logistic function with linear scaling and displacement used in Experiment 2. The weights were chosen as follows: $\mathbf{w}_1(k) = 1 + k/100$, $\mathbf{w}_2(k) = \sin(0.06k) - 2$, $\mathbf{w}_3(k) = 0.1$, $\mathbf{w}_4(k) = 1$, $\mathbf{w}_5(k) = -0.5$.

weights and generated another 200 input-output data test sets from the original model. The input test data was then fed to the trained model, using the weights values estimated at the 200-th time step. Subsequently, the output prediction of the trained model was compared to the output data from the original model so as to assess the generalisation performance of the training process.

We repeated the experiment using two different values of Q, as shown in Figures 11 and 12. With a very small value of Q (0.0001), the algorithm was able to estimate the time-varying latent weights. After 100 time steps, the algorithm became very confident, as indicated by the narrow error bars. By increasing Q to 0.01, we found that the algorithm becomes less confident. Yet, it is able to explore wider regions of parameter space as shown in the histogram of Figure 13.

The variance (Q) of the noise added in the prediction stage implies a trade-off between the algorithms' confidence and the size of the search region in parameter space. Too large a value of Q will allow the algorithm to explore a large region of the space, but will never allow the algorithm to converge to specific error minimum. A very small value, on the other hand, will cause all the sampling trajectories to converge to a very narrow region. The problem with the latter is that if the process generating the data changes with time, we want to keep the search region large enough so as to find the new minima quickly and efficiently. To address this problem we have employed a very simple heuristic. Specifically, when we add Gaussian noise in the prediction stage, we select a small set of the trajectories randomly and then add to them Gaussian noise with a covariance much larger than Q. In theory, this should allow the algorithm to converge and yet allow it to keep exploring other regions of parameter space. From an evolutionary computing perspective, we can think of these few randomly selected trajectories as mutations. The results of applying this idea to our problem are shown in Figure 13. They suggest that our heuristic reduces the trade-off between the algorithms' confidence and the size of the search region. Yet, we believe that further research on algorithms for adapting Q is necessary.

8.3 Example with Real Data

An interesting source of inference problems requiring nonlinear modelling in a time-varying environment is financial data analysis. Many applications ranging from time series prediction to stock selection have been attempted using neural networks. Hutchinson, Lo and Poggio (Hutchinson, Lo and Poggio 1994) and Niranjan (Niranjan 1996) look at the problem of pricing options contracts in a data driven manner. The relationship between the price of an options contract and the parameters relating to the option, such as the price of the underlying asset and the time to maturity, is known to be nonlinear. Hutchinson *et. al.* show that a neural network can form a very good approximation to the price of the options contract as a function of variables relating to this. This relationship is also likely to change over the lifetime of the option. Hence one would expect that a sequential approach for tracking the price of the options contract might be a more natural way to handle the non-stationarity. Niranjan addresses this by means of an extended Kalman filter. We extend on the above work, using the data used in (Niranjan 1996) consisting of daily FTSE100 options during the Feb 93 - Dec 93 period. We compare the approximation and convergence capabilities of a number of algorithms discussed in earlier sections. The models considered are:

Trivial : This method simply involves using the current value of the option as the next prediction.

RBF-EKF : Represents a regularised radial basis function network with 4 hidden neurons, which was originally proposed in (Hutchinson et al. 1994). The output weights are estimated with a Kalman filter, while the means of the radial functions correspond to random subsets of the data and their covariance is set to the identity matrix as in (Niranjan 1996).

BS : Corresponds to a conventional Black-Scholes model with two outputs (normalised call and put prices) and two parameters (risk-free interest rate and volatility). The risk-free interest rate was set to 0.06, while the volatility was estimated over a moving window (of 50 time steps) as described in (Hull 1997).

MLP-EKF : A multi-layer perceptron, with 6 sigmoidal hidden units and a linear output neuron, trained with the EKF algorithm (de Freitas et al. 1997). The noise covariances R and Q and the states covariance P were set to diagonal matrices with entries equal to 10^{-6}, 10^{-7} and 10 respectively. The weights prior corresponded to a zero mean Gaussian density with covariance equal to 1.

MLP-EKFE : Represents an MLP, with 6 sigmoidal hidden units and a linear output neuron, trained with a hierarchical Bayesian model, whereby the weights are estimated

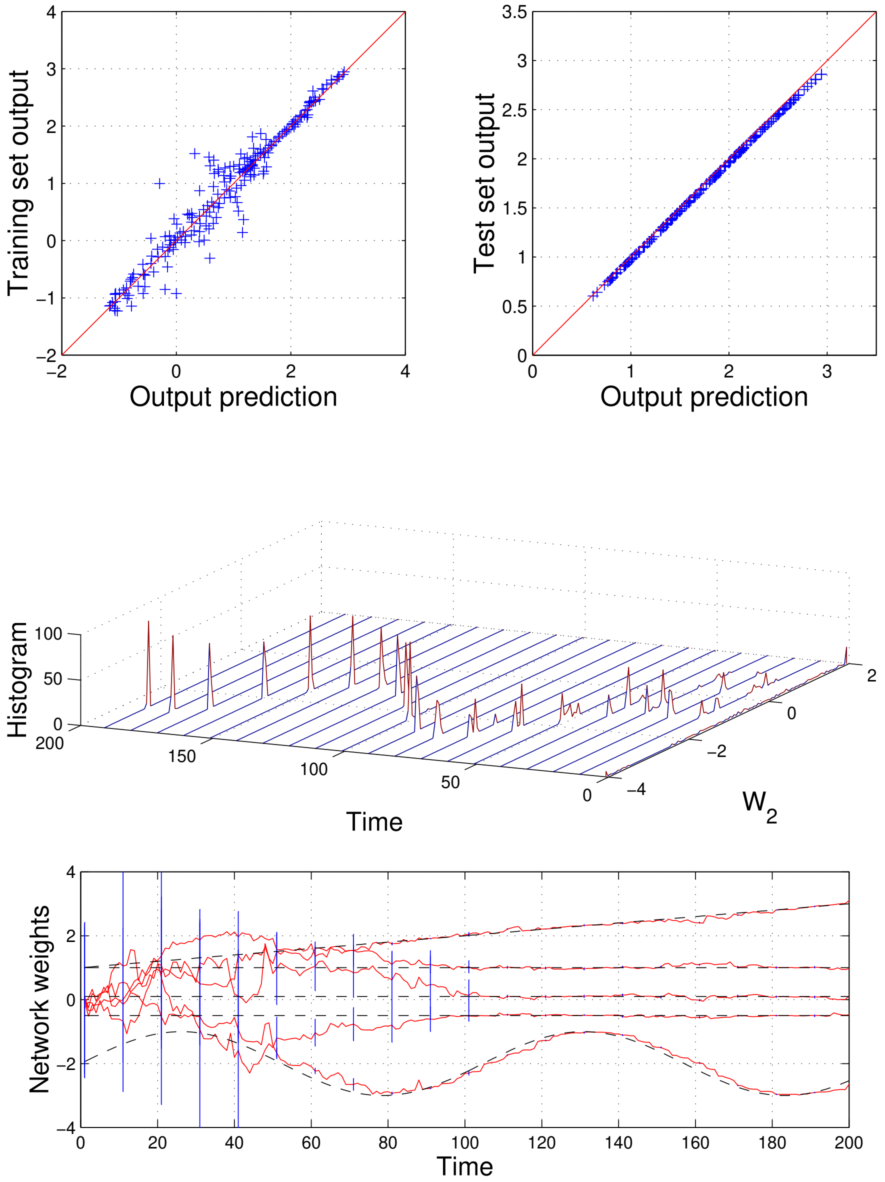

Figure 11: Weights tracking with $Q = 0.0001$. The top left plot shows the one-step-ahead prediction errors in the training set. The top right plots shows the prediction errors on validation data, assuming that after the 200th time step the weights stop varying with time. The middle plot shows the histogram of $\mathbf{w}_2$. Finally, the bottom plot shows the tracking of the weights with the posterior mean estimate computed by the HySIR. The bottom plot also shows the one standard deviation error bars of the estimator. For this particular value of Q, the estimator becomes very confident after 100 iterations and, consequently, the histogram for $\mathbf{w}_2$ converges to a narrow Gaussian distribution.

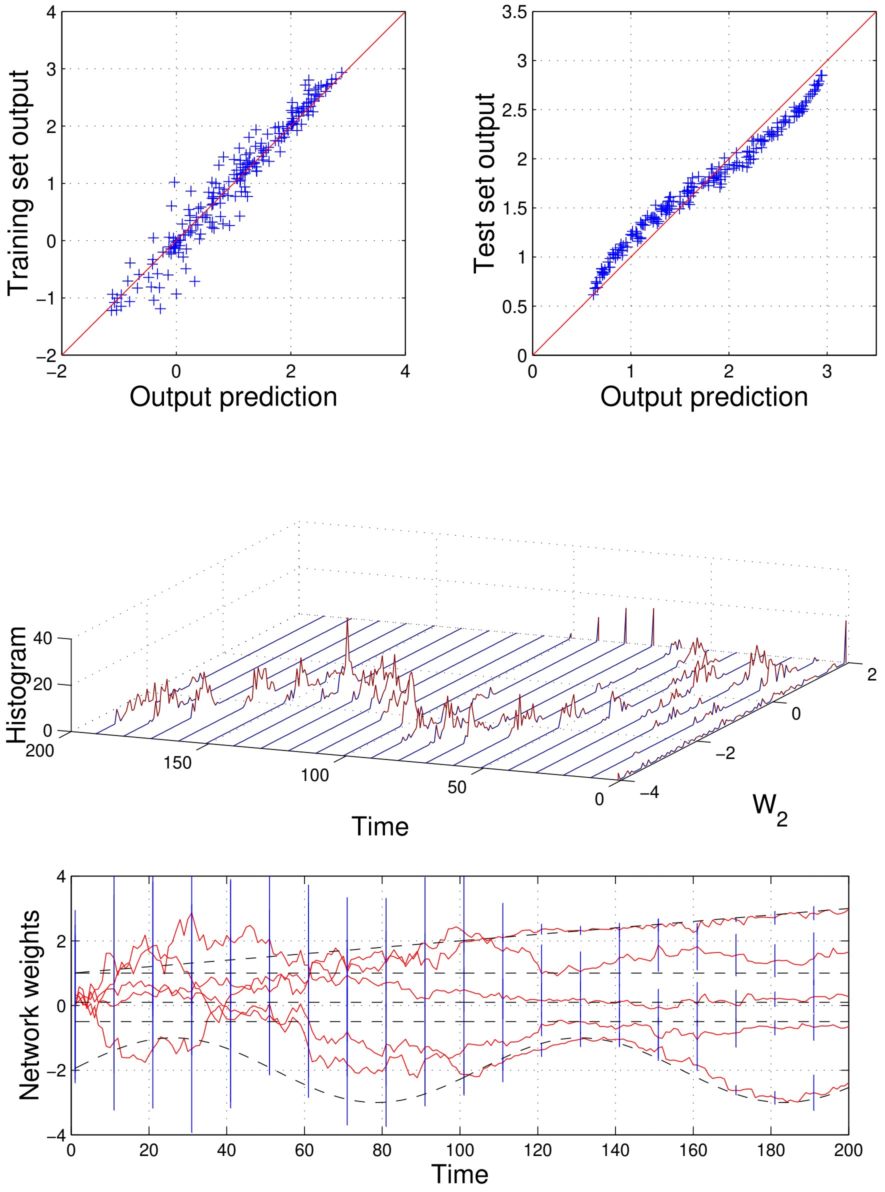

Figure 12: Weights tracking with $Q = 0.01$. The top left plot shows the one-step-ahead prediction errors in the training set. The top right plots shows the prediction errors on validation data, assuming that after the 200th time step the weights stop varying with time. The middle plot shows the histogram of $\mathbf{w}_2$. Finally, the bottom plot shows the tracking of the weights with the posterior mean estimate computed by the HySIR. The bottom plot also shows the one standard deviation error bars of the estimator. For this value of Q, the search region is particularly large (wide histograms), but the algorithm lacks confidence (wide error bars).

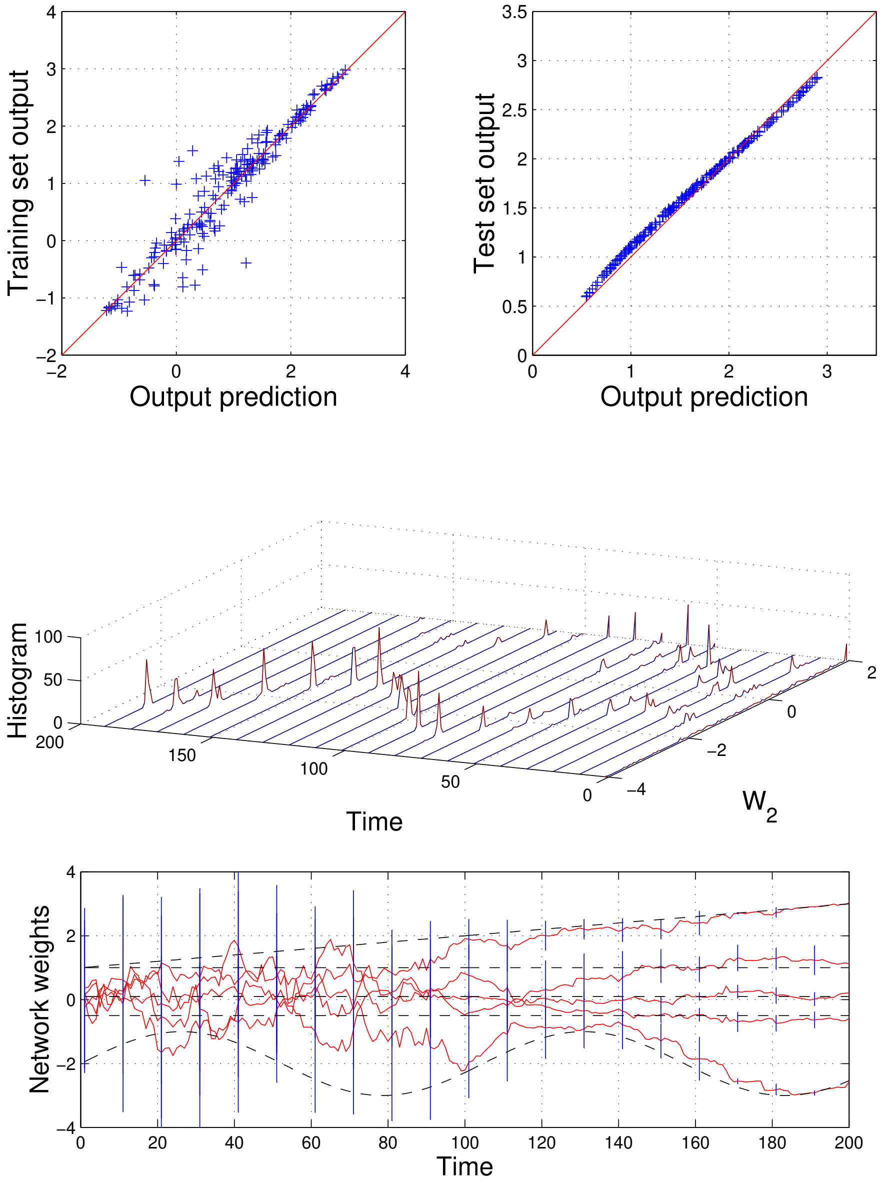

Figure 13: Weights tracking with $Q = 0.0001$ and two random mutations per time step with $Q = 0.1$. The top left plot shows the one-step-ahead prediction errors in the training set. The top right plots shows the prediction errors on validation data, assuming that after the 200th time step the weights stop varying with time. The middle plot shows the histogram of $\mathbf{w}_2$. Finally, the bottom plot shows the tracking of the weights with the posterior mean estimate computed by the HySIR. The bottom plot also shows the one standard deviation error bars of the estimator. The histogram shows that the mutations expand the search region considerably.

with the EKF algorithm and the noise statistics are computed by maximising the evidence density function (de Freitas et al. 1997). The states covariance P was given by a diagonal matrix with entries equal to 10. The weights prior corresponded to a zero mean Gaussian density with covariance equal to 1.

MLP-SIR : Corresponds to an MLP, with 6 sigmoidal hidden units and a linear output neuron, trained with the SIR algorithm. 1000 samples were employed in each simulation. The noise covariances R and Q were set to diagonal matrices with entries equal to 10^{-4} and 10^{-5} respectively. The prior samples for each layer were drawn from a zero mean Gaussian density with covariance equal to 1.

MLP-HySIR : Corresponds to an MLP, with 6 sigmoidal hidden units and a linear output neuron, trained with the Kalman HySIR algorithm. Each simulation made use of 10 samples. The noise covariances R and Q were set to diagonal matrices with entries equal to 1 and 10^{-6} respectively. The prior samples for each layer were drawn from a zero mean Gaussian density with covariance equal to 1. The Kalman filter parameters R^*, Q^* and P were set to diagonal matrices with entries equal to 10^{-4}, 10^{-5} and 1000 respectively.

Strike price	2925	3025	3125	3225	3325
Trivial	0.0783	0.0611	0.0524	0.0339	0.0205
RBF-EKF	0.0538	0.0445	0.0546	0.0360	0.0206
BS	0.0761	0.0598	0.0534	0.0377	0.0262
MLP-EKF	0.0414	0.0384	0.0427	0.0285	0.0145
MLP-EKFE	0.0404	0.0366	0.0394	0.0283	0.0150
MLP-SIR	0.0419	0.0405	0.0432	0.0314	0.0164
MLP-HySIR	0.0394	0.0362	0.0404	0.0273	0.0151

Table 2: One-step-ahead prediction errors on call options.

Table 2 shows the mean squared error between the actual value of the options contracts and the one-step-ahead prediction from the various algorithms. Though the differences are small, one sees that sequential algorithms produce lower prediction errors, and that the HySIR algorithm produces even lower errors than the Kalman filter. These errors are computed on data corresponding to the last 100 days, to allow for convergence of the algorithms. Note that in the on-line setting considered here, all these errors are computed on unseen data.

Figure 14 shows the approximations of the MLP-HySIR algorithm on the call and put option prices at the same strike price (3325). The algorithm exhibits fast convergence and accurate tracking. In Figure 15, we plot the probability density function over the network's one-step-ahead prediction, computed through 100 samples propagated by the HySIR approach.

Typically, the one-step-ahead predictions for a group of options on the same cash product, but with different strike prices or duration to maturity, can be used to determine whether one of the options is being mispriced. Knowing the probability distribution of the network outputs allows us to design more interesting pricing tools.

9 Conclusions

In this paper, we have presented a sequential Monte Carlo approach for training neural networks in a Bayesian setting. This approach enables accurate characterisation of the probability distributions which tend to be very complicated for nonlinear models such as neural

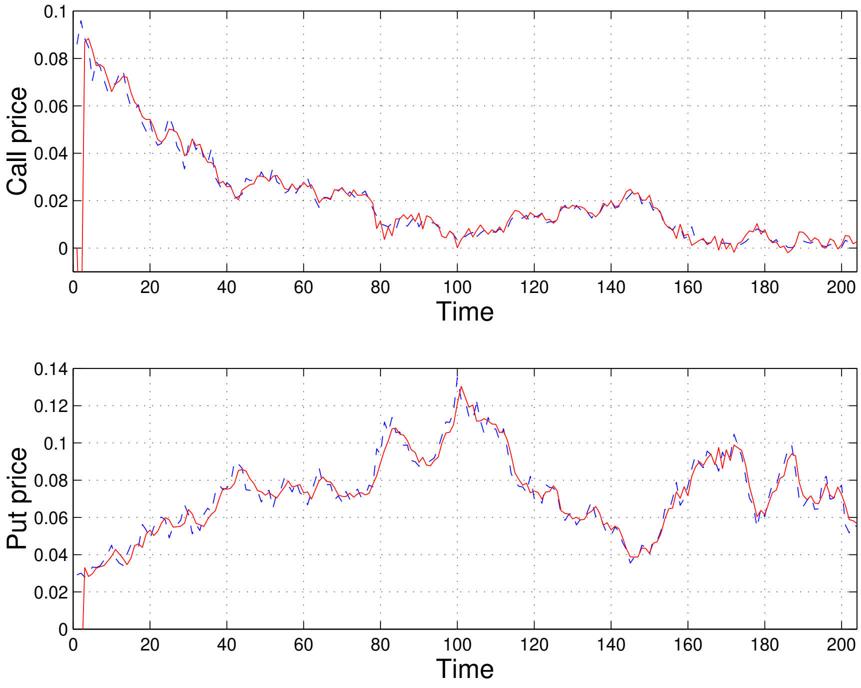

Figure 14: Tracking call and put option prices with the MLP-HSIR method. Estimated values [—] and actual values [- -].

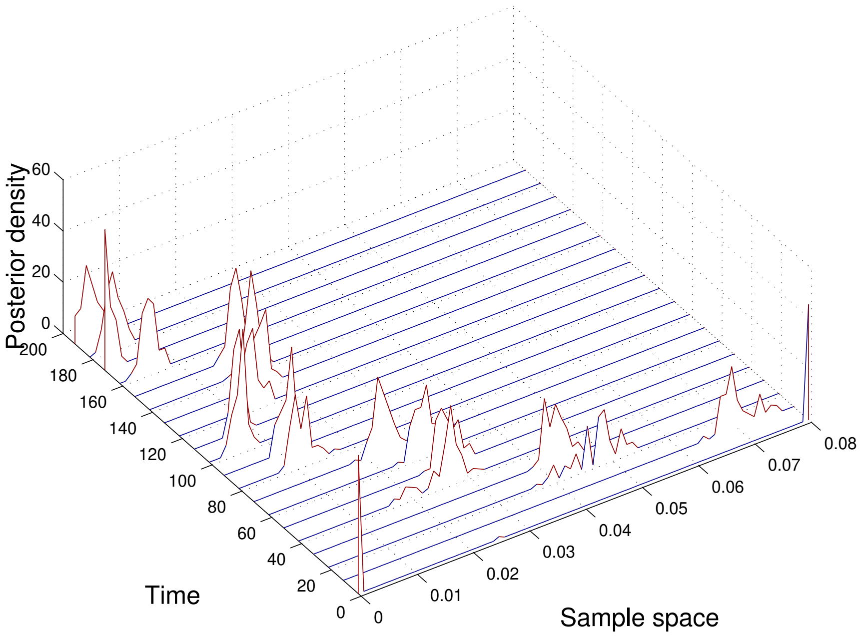

Figure 15: Probability density function of the neural networks' one-step-ahead prediction of the call price for the FTSE option with a strike price of 3325.

networks. Propagating a number of samples in a state space dynamical systems formulation of the problem is the key idea in this framework.

Experimental illustrations presented here suggest that it is possible to effectively apply sampling methods to tackle difficult problems involving nonlinear and time varying processes, typical of many sequential signal processing problems.

A particular contribution in this paper is the HySIR algorithm, which combines gradient information of the error function with probabilistic sampling information. Simulations show that this combination performs better than EKF and conventional SIR approaches. HySIR can be interpreted as a Gaussian mixture filter, in that only a few sampling trajectories need to be employed. Yet, as the number of trajectories increases, the computational requirements increase only linearly. Therefore, the method is also suitable as a sampling strategy for approximating multi-modal distributions.

Our experiment with financial data showed that the sequential sampling approach, in addition to matching the one-step-ahead square errors of the best available pricing methods, allowed us to obtain complete estimates of the probability density functions of the one-step-ahead predictions. Further avenues of research include the design of algorithms for adapting the noise covariances R and Q, studying the effect of different noise models for the network weights and improving the computational efficiency of the algorithms.

10 Acknowledgements

We would like to thank Neil Gordon (DERA, UK), Sue Johnson (Department of Engineering, Cambridge) and Michael Isard (Computer Vision Group, Oxford) for interesting and useful discussions on this work. João FG de Freitas is financially supported by two University of the Witwatersrand Merit Scholarships, a Foundation for Research Development Scholarship (South Africa), an ORS award and a Trinity College External Studentship (Cambridge).

References

Akashi, H. and Kumamoto, H. (1977). Random sampling approach to state estimation in switching environments, *Automatica* **13**: 429–434.

Avitzour, D. (1995). A stochastic simulation Bayesian approach to multitarget tracking, *IEE Proceedings on Radar, Sonar and Navigation* **142**(2): 41–44.

Bar-Shalom, Y. and Li, X. R. (1993). *Estimation and Tracking: Principles, Techniques and Software*, Artech House, Boston.

Beadle, E. R. and Djuric, P. M. (1997). A fast weighted Bayesian bootstrap filter for non-linear model state estimation, *IEEE Transactions on Aerospace and Electronic Systems* **33**(1): 338–343.

Berzuini, C., Best, N. G., Gilks, W. R. and Larizza, C. (1997). Dynamic conditional independence models and Markov Chain Monte Carlo methods, *Journal of the American Statistical Association* **92**(440): 1403–1412.

Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford.

Brass, A., Pendleton, B. J., Chen, Y. and Robson, B. (1993). Hybrid Monte Carlo simulations theory and initial comparison with molecular dynamics, *Biopolymers* **33**(8): 1307–1315.

Carpenter, J., Clifford, P. and Fearnhead, P. (1997). An improved particle filter for non-linear problems, *Technical report*, Department of Statistics, Oxford University, England. Available at http://www.stats.ox.ac.uk/~clifford/index.htm.

Collings, I. B., Krishnamurthy, V. and Moore, J. B. (1994). On-line identification of hidden Markov models via recursive prediction error techniques, *IEEE Transactions on Signal Processing* **42**(12): 3535–3539.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function, *Mathematics of Control, Signals, and Systems* **2**(4): 303–314.

de Freitas, J. F. G., Niranjan, M. and Gee, A. H. (1997). Hierarchichal Bayesian-Kalman models for regularisation and ARD in sequential learning, *Technical Report CUED/F-INFENG/TR 307*, Cambridge University, http://svr-www.eng.cam.ac.uk/~jfgf.

de Freitas, J. F. G., Niranjan, M. and Gee, A. H. (1998a). The EM algorithm and neural networks for nonlinear state space estimation, *Technical Report CUED/F-INFENG/TR 313*, Cambridge University, http://svr-www.eng.cam.ac.uk/~jfgf.

de Freitas, J. F. G., Niranjan, M. and Gee, A. H. (1998b). Nonlinear state space learning with EM and neural networks, *IEEE International Workshop on Neural Networks in Signal Processing*, Cambridge, England.

de Freitas, J. F. G., Niranjan, M. and Gee, A. H. (1998c). Regularisation in sequential learning algorithms, *in* M. I. Jordan, M. J. Kearns and S. A. Solla (eds), *Advances in Neural Information Processing Systems*, Vol. 10, MIT Press.

Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society Series B* **39**: 1–38.

Doucet, A. (1997). *Monte Carlo Methods for Bayesian estimation of hidden Markov Models. Application to Radiation Signals*, PhD thesis, University Paris-Sud, Orsay, France. Chapters 4 and 5 in English.

Doucet, A. (1998). On sequential simulation-based methods for Bayesian filtering, *Technical Report CUED/F-INFENG/TR 310*, Cambridge University. Available at http://www.stats.bris.ac.uk:81/MCMC/pages/list.html.

Duane, S., Kennedy, A. D., Pendleton, B. J. and Roweth, D. (1987). Hybrid Monte Carlo, *Physics Letters B* **195**(2): 216–222.

Efron, B. (1982). *The Bootstrap, Jacknife and other Resampling Plans*, CBMS-NSF Regional Conference Series in Applied Mathematics 38, SIAM, Philadelphia.

Elliott, R. J. (1994). Exact adaptive filters for Markov chains observed in Gaussian noise, *Automatica* **30**(9): 1399–1408.

Gelfand, A. E., Dey, D. K. and Chang, H. (1992). Model determination using predictive distributions with implementation via sampling-based methods, *in* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 4*, pp. 147–167.

Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. (1995). *Bayesian Data Analysis*, Chapman and Hall.

Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration, *Econometrica* **24**: 1317–1399.

Gilks, W. R., Richardson, S. and Spiegelhalter, D. J. (eds) (1996). *Markov Chain Monte Carlo in Practice*, Chapman and Hall, Suffolk.

Gordon, N. and Whitby, A. (1995). A Bayesian approach to target tracking in the presence of a glint, *in* O. E. Drummond (ed.), *Signal and Data Processing of Small Targets*, Vol. SPIE 2561, pp. 472–483.

Gordon, N. J. (1994). *Bayesian Methods for Tracking*, PhD thesis, Imperial College, University of London.

Gordon, N. J., Salmond, D. J. and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation, *IEE Proceedings-F* **140**(2): 107–113.

Hammersley, J. H. and Handscomb, D. C. (1968). *Monte Carlo Methods*, Methuen, London.

Handschin, J. E. (1970). Monte Carlo techniques for prediction and filtering of non-linear stochastic processes, *Automatica* **6**: 555–563.

Handschin, J. E. and Mayne, D. Q. (1969). Monte Carlo techniques to estimate the conditional expectation in multi-stage non-linear filtering, *International Journal of Control* **9**(5): 547–559.

Harvey, A. C. (1970). *Forecasting, Structural Time Series Models, and the Kalman Filter*, Cambridge University Press.

Hull, J. C. (1997). *Options, Futures, and Other Derivative*, third edn, Prentice Hall.

Hutchinson, J. M., Lo, A. W. and Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks, *The Journal of Finance* **49**(3): 851–889.

Isard, M. and Blake, A. (1996). Contour tracking by stochastic propagation of conditional density, *European Conference on Computer Vision*, Cambridge, UK, pp. 343–356.

Jazwinski, A. H. (1969). Adaptive filtering, *Automatica* **5**: 475–485.

Jazwinski, A. H. (1970). *Stochastic Processes and Filtering Theory*, Academic Press.

Kadirkamanathan, V. and Kadirkamanathan, M. (1995). Recursive estimation of dynamic modular RBF networks, *in* D. S. Touretzky, M. C. Mozer and M. E. Hasselmo (eds), *Advances in Neural Information Processing Systems 8*, pp. 239–245.

Kitagawa, G. (1987). Non-Gaussian state-space modeling of nonstationary time series, *Journal of the American Statistical Association* **82**(400): 1032–1063.

Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models, *Journal of Computational and Graphical Statistics* **5**: 1–25.

Kong, A., Liu, J. S. and Wong, W. H. (1994). Sequential imputations and Bayesian missing data problems, *Journal of the American Statistical Association* **89**(425): 278–288.

Kramer, S. C. and Sorenson, H. W. (1988). Recursive Bayesian estimation using piece-wise constant approximations, *Automatica* **24**(6): 789–801.

Krishnamurthy, V. and Moore, J. B. (1993). On-line estimation of hidden Markov model parameters based on the Kullback-Leibler information measure, *IEEE Transactions on Signal Processing* **41**(8): 2557–2573.

Le Cun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E. and Hubbard, W. (1989). Backpropagation applied to handwritten zip code recognition, *Neural Computation* **1**: 541–551.

Li, X. R. and Bar-Shalom, Y. (1994). A recursive multiple model approach to noise identification, *IEEE Transactions on Aerospace and Electronic Systems* **30**(3): 671–684.

Liu, J. S. and Chen, R. (1995). Blind deconvolution via sequential imputations, *Journal of the American Statistical Association* **90**(430): 567–576.

Liu, J. S. and Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems, To appear in the Journal of the American Statistical Association.

Mackay, D. J. C. (1992). Bayesian interpolation, *Neural Computation* **4**(3): 415–447.

Müller, P. (1991). Monte Carlo integration in general dynamic models, *Contemporary Mathematics* **115**: 145–163.

Müller, P. (1992). Posterior integration in dynamic models, *Computer Science and Statistics* **24**: 318–324.

Müller, P. and Rios Insua, D. (1998). Issues in Bayesian analysis of neural network models, *Neural Computation* **10**: 571–592.

Neal, R. M. (1996). *Bayesian Learning for Neural Networks*, Lecture Notes in Statistics No. 118, Springer-Verlag, New York.

Niranjan, M. (1996). Sequential tracking in pricing financial options using model based and neural network approaches, *in* M. C. Mozer, M. I. Jordan and T. Petsche (eds), *Advances in Neural Information Processing Systems*, pp. 960–966.

Pitt, M. K. and Shephard, N. (1997). Filtering via simulation: Auxiliary particle filters, *Technical report*, Department of Statistics, Imperial College of London, England. Available at http://www.nuff.ox.ac.uk/economics/papers.

Pole, A. and West, M. (1990). Efficient Bayesian learning in non-linear dynamic models, *Journal of Forecasting* **9**: 119–136.

Puskorius, G. V. and Feldkamp, L. A. (1991). Decoupled extended Kalman filter training of feedforward layered networks, *International Joint Conference on Neural Networks*, Seattle, pp. 307–312.

Puskorius, G. V., Feldkamp, L. A. and Davis, L. I. (1996). Dynamic neural network methods applied to on-vehicle idle speed control, *Proceedings of the IEEE* **84**(10): 1407–1420.

Rauch, H. E., Tung, F. and Striebel, C. T. (1965). Maximum likelihood estimates of linear dynamic systems, *AIAA Journal* **3**(8): 1445–1450.

Robinson, T. (1994). The application of recurrent nets to phone probability estimation, *IEEE Transactions on Neural Networks* **5**(2): 298–305.

Rubin, D. B. (1988). Using the SIR algorithm to simulate posterior distributions, *in* J. M. Bernardo, M. H. DeGroot, D. V. Lindley and A. F. M. Smith (eds), *Bayesian Statistics 3*, pp. 395–402.

Ruck, D. W., Rogers, S. K., Kabrisky, M., Maybeck, P. S. and Oxley, M. E. (1992). Comparative analysis of backpropagation and the extended Kalman filter for training multilayer perceptrons, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(6): 686–690.

Shah, S., Palmieri, F. and Datum, M. (1992). Optimal filtering algorithms for fast learning in feedforward neural networks, *Neural Networks* **5**: 779–787.

Sibisi, S. (1989). Regularization and inverse problems, *Maximum Entropy and Bayesian Methods*, Kluwer Academic Publishers, pp. 389–396.

Singhal, S. and Wu, L. (1988). Training multilayer perceptrons with the extended Kalman algorithm, *in* D. S. Touretzky (ed.), *Advances in Neural Information Processing Systems*, Vol. 1, San Mateo, CA, pp. 133–140.

Smith, A. F. M. and Gelfand, A. E. (1992). Bayesian statistics without tears: a sampling-resampling perspective, *American Statistician* **46**: 84–88.

Sorenson, H. W. and Alspach, D. L. (1971). Recursive Bayesian estimation using Gaussian sums, *Automatica* **7**: 465–479.

Zaritskii, V. S., Svetnik, V. B. and Shimelevich, L. I. (1975). Monte-Carlo techniques in problems of optimal information processing, *Automatic Remote Control* pp. 2015–2022.