Online Learning with Radial Basis Function Networks

Gabriel Borrageiro

GABRIEL.BORRAGEIRO.20@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Nick Firoozye

N.FIROOZY@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Paolo Barucca

P.BARUCCA@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Abstract

We investigate the benefits of feature selection, nonlinear modelling and online learning when forecasting in financial time series. We consider the sequential and continual learning sub-genres of online learning. The experiments we conduct show that there is a benefit to online transfer learning, in the form of radial basis function networks, beyond the sequential updating of recursive least-squares models. We show that the radial basis function networks, which make use of clustering algorithms to construct a kernel Gram matrix, are more beneficial than treating each training vector as separate basis functions, as occurs with kernel Ridge regression. We demonstrate quantitative procedures to determine the very structure of the radial basis function networks. Finally, we conduct experiments on the log returns of financial time series and show that the online learning models, particularly the radial basis function networks, are able to outperform a random walk baseline, whereas the offline learning models struggle to do so.

Keywords: online learning, transfer learning, radial basis function networks, financial time series, multi-step forecasting

1. Introduction

Financial time series are characterised by high serial autocorrelation and nonstationarity. The classic paper on the theory of option pricing by Black and Scholes (1973), assumes that stock prices follow a geometric Brownian motion through time, which produces a log-normal distribution of price returns. Merton (1976) modelled the dynamics of financial assets as a jump-diffusion process, which is now commonly used in financial econometrics. The jump-diffusion process implies that financial time series should observe small changes over time, so-called continuous changes, as well as occasional jumps. One approach for coping with nonstationarity, is to continuously learn online. Online learning can be classified into three broad areas: sequential updating, states of nature or transitional learning, and continual learning.

Sequential learning in time is a self-explanatory concept and has a rich history. The Kalman filter (Kalman, 1960) is a state space model which was originally designed for tracking objects in time, such as airplanes or missiles, from noisy measurements, such

as radar. Several approaches exist for sequential learning in nonstationary data. These include discounted least squares (Abraham and Ledolter, 1983) and exponentially weighted recursive least squares (Liu et al., 2010). Barber et al. (2011) consider Bayesian approaches to time series modelling which are amenable to sequential learning.

Transitional learning includes reinforcement learning, which allows agents to interact with their environment, mapping situations to actions so as to maximise a numerical award. Well known online reinforcement learning algorithms include q-learning (Watkins, 1989) and sarsa (Rummery and Niranjan, 1994). The classical k-armed bandit problem, where one is faced with a choice amongst k possible options, is formulated as an online learning problem. After each choice, a reward is assigned. The interplay between hedonistic exploitation and potentially costly exploration, leads to a number of algorithms such as ϵ -greedy, stochastic gradient ascent and upper confidence bound bandits (Sutton and Barto, 2018).

Continual learning is an area of study that asks how artificial systems might learn sequentially, as biological systems do, from a continuous stream of correlated data (Hadsell et al., 2020). They include gradient based methods (Kaplanis et al. 2018, Kirkpatrick et al. 2017), meta learning (Wang et al., 2017) and transfer learning (Yang et al., 2020). The goal of transfer learning is broadly to transfer knowledge from one model (the source) to another (the target). Sub-paradigms of transfer learning include inductive, where labelled data is available in the target domain, and transductive, where labelled data is available only in the source domain. Koshiyama et al. (2020) applies transfer learning to systematic trading strategy development, with goals of minimising backtest overfitting and generating higher risk adjusted returns.

In our paper, we aim to assess the benefits of predictive modelling, feature selection, online learning and nonlinear modelling. We limit the scope of our experimentation to financial time series, which at times exhibit high autocorrelation, nonstationarity and non-linearity. This behaviour might be classified as concept drift (Iwashita and Papa, 2019). On its own, the sequential learning approach to online learning might be suboptimal to an approach that includes other forms of online learning such as transfer learning. The marquis model of this paper, an online learning radial basis function network, employs two models. The unsupervised learning step, which learns the hidden processing units of the network, is capable of time series pattern recognition. The supervised learning step, which maps these hidden processing units to the output, employs sequential learning. The model thus straddles the continual learning and sequential learning domains. For this reason, we explore its benefits within forecasting financial time series.

Through experimentation, we wish to ask several research questions and test several hypotheses. A first question would be, should predictive modelling consistently outperform model free approaches as it pertains to forecast accuracy? A second research question is, is there a benefit to doing feature selection? Related to this question, should we look at exogenous features? The likes of Chatfield (2019) caution against the indiscriminate use of exogenous features. The reasons he gives include the introduction of serial autocorrelation in the regression residuals, leading to spurious regressions (Granger and Newbold, 1974) and the belief that exogenous features should provide clear contextual reasons as to why they should be included, rather than considering optimisation criteria alone, such as squared-error loss minimisation. A further research question is, even after making time series stationary, should online learning outperform offline learning? The hypothesis of a jump-diffusion process implies that even if returns are constructed, and tests such as the Dickey and Fuller (1979) imply data stationarity, occasional jumps might still occur, which render the offline learning models less useful. We expect a priori for the relationship between a set of predictors and a target to be changing over time, thus requiring some form of sequential learning. To test this hypothesis, we create some online learning models and measure their

performance against some offline learning models. All models are then baselined against a random walk model.

1.1 Our Contribution

Our paper contributes to the research of continual learning in financial time series. This is important as our experiment shows that continual learning provides a benefit with multi-step forecasting, above and beyond sequential learning. If we compare the local learning of the radial basis function network with the global learning technique of the feed-forward neural network, the latter suffers from catastrophic forgetting. The radial basis function network that we formulate, is naturally designed to measure the similarity between test samples and continuously updated prototypes that capture the characteristics of the feature space. As such, the model is somewhat robust in mitigating catastrophic forgetting.

Although closely related to kernel Ridge regression, our experiment results show that the radial basis function networks, which make use of clustering algorithms, and therefore have condensed hidden processing units, are more beneficial than treating each training vector as separate basis functions. Indeed, considering the regret bounds of section 4, then the regret with respect to the expert with forward looking bias, scales with the number of experts. Thus, care should be taken as to how those experts are formulated and chosen. Furthermore, if one makes use of the Nyström approximation to the kernel Gram matrix, as shown in algorithm 2, random sub-sampling of the feature space is of less benefit than employing clustering or pseudo-clustering algorithms to construct a kernel Gram matrix for use in modelling non-stationary data, such as financial time series.

In terms of selecting the structure of the radial basis function network itself, namely selection of the hidden processing units, we have demonstrated quantitative procedures to do so, which are less expensive than procedures based on cross-validation or generalised cross-validation alone. For the radial basis function network that makes use of kmeans++, we consider several possible cluster structures and select the one that maximises the silhouette score, see section 3.2.1. For the radial basis function network that makes use of Gaussian mixture models, we use the meta-algorithm and modified expectation-maximisation procedure outlined by Figueiredo and Jain (2002) and discussed in section 3.2.2, to select mixtures that provide both mixture selection and model estimation. In addition, the specific formulation annihilates mixtures that are not supported by the data.

Finally, we show a feature selection meta-algorithm that combines two algorithms, namely forward stepwise selection and variance inflation factor minimisation. The former is used to select features that have explanatory power with respect to the response, even in a high dimensional feature space setting, and the latter is used to prune any correlated features back, which are likely to result in greater prediction variance.

2. Preliminaries

We begin with a discussion of the baseline model and the preliminary models that we use in our experimentation. Also discussed is a feature selection meta-algorithm that we make use of.

2.1 The Random Walk Model

As described in Harvey (1993), the simplest non-stationary process is the random walk model

$$y_t = y_{t-1} + \epsilon_t.$$

This is the model we use as a baseline in the experiments we conduct. The first differences of this model, $y_t - y_{t-1}$, are stationary, which leads to a general class of models known as autoregressive integrated moving average (arima). Repeatedly substituting for past values gives

$$y_t = y_0 + \sum_{i=1}^{t-1} \epsilon_{t-i}.$$

The expectation of the random walk model is

$$\mathbb{E}[y_t] = y_0,$$

indicating a constant mean over time. The variance of a random walk process

$$\text{Var}[y_t] = t\sigma^2,$$

and covariance

$$\text{Cov}[y_t, y_{t-\tau}] = |t - \tau|\sigma^2,$$

is however, non-stationary. It is common within econometrics to take log first differences, as this has a variance stabilising effect. Let us denote the h -step ahead log return as

$$y_{t+h} = \log(Y_{t+h}/Y_t) = \sum_{i=t+1}^h \log(Y_i/Y_{i-1}),$$

where, in the context of financial time series for example, Y_t would typically represent a quoted price at time t . When forecasting using the random walk model in log returns space, the mean square prediction error (mspe) for the h' th forecast horizon becomes

$$\text{mspe}_h = \frac{1}{t-h} \sum_{i=1}^{t-h} (y_{i+h} - \hat{y}_0)^2 = \frac{1}{t-h} \sum_{i=1}^{t-h} y_{i+h}^2.$$

A large amount of academic literature shows that it is difficult to beat the random walk model when forecasting returns of financial time series. For example, Meese and Rogoff (1983) show that the monetary, econometric model is unable to outperform a random walk model when forecasting currency exchange rates, which implies that exchange rates behave in a purely random and unpredictable manner. This phenomenon is known as the Meese–Rogoff puzzle. Engel (1994) worked further on this puzzle. He fits a Markov-switching model to 18 exchange rates at sampled quarterly frequencies, and finds that the model fits well in-sample for many exchange rates. However, by the mean square prediction error criterion, the Markov model does not generate superior forecasts to a random walk for the forward exchange rate. He finds some evidence to support that the forecasts of the Markov-switching model are superior at predicting the direction of change of the exchange rate. More recently, Ince et al. (2019) provide exchange rate forecasting by combining vector autoregression models with multi-layer feed-forward neural networks. Their models are able to outperform the random walk baseline on currency returns sampled monthly and forecasted one step ahead. They conclude that employing alternative artificial neural network structures such as radial basis function networks and recurrent neural networks remain as relevant future research topics.

2.2 The Autoregressive(1) Model

As Harvey (1993) discusses, the autoregressive model of order 1

$$y_t = \phi_1 y_{t-1} + \epsilon_t, \quad t = 1, \dots, T,$$

is closely related to the random walk model. We denote this model as $\text{ar}(1)$. Substituting repeatedly for lagged values of y_t gives

$$y_t = \sum_{i=1}^{t-1} \phi^i \epsilon_{t-i} + \phi^t y_0.$$

The expectation of the $\text{ar}(1)$ model is

$$\mathbb{E}[y_t] = E \left[\sum_{i=1}^{t-1} \phi^i \epsilon_{t-i} \right] + \mathbb{E}[\phi^t y_0] = \phi^t y_0.$$

For $|\theta| \geq 1$, the mean value of the process depends on the starting value, y_0 . For $|\theta| < 1$, the impact of the starting value is negligible asymptotically. The variance of the $\text{ar}(1)$ process is

$$\gamma[0] = E \left[\sum_{i=1}^{t-1} \phi^i \epsilon_{t-i} \right]^2 = \sigma^2 \sum_{i=1}^{\infty} \phi^{2i} = \sigma^2 / (1 - \phi^2).$$

2.3 Feature Selection

Linear regression is a model of the form

$$p(y|\mathbf{x}, \boldsymbol{\theta}) = \mathcal{N}(y|\mathbf{w}^T \mathbf{x}, \sigma^2),$$

where y is the response, $\mathbf{x}$ is the vector of independent variables, and the parameter set $\boldsymbol{\theta} = [\mathbf{w}, \sigma^2]$, must be estimated. The assumption is that the conditional probability $p(y|\mathbf{x}, \boldsymbol{\theta})$, is normally distributed. How should we select the predictor set $\mathbf{x}$, where $\mathbf{x} \subset \mathcal{X}$? In the spirit of Occam's razor, we seek a subset of minimally correlated predictors, with maximal explanatory power. There are various feature selection algorithms, which Hastie et al. (2009) discuss in detail. Here, we show a feature selection meta-algorithm that combines two algorithms, namely forward stepwise selection and variance inflation factor minimisation. The former is used to select features that have explanatory power with respect to the response, even in a high dimensional feature space setting, and the latter is used to prune any correlated features back. Correlated features are likely to result in greater prediction variance. In the algorithm that follows, κ is the maximum variance

inflation factor that is permitted, $\mathbf{X}$ is the $n \times d$ predictor matrix, and $\mathbf{y}$ is the $n \times 1$ target vector.

Algorithm 1: forward stepwise selection with variance inflation factors

```

Require:  $\kappa$ 
Initialise:  $\mathbb{S} = [1, 2, \dots, d]$ ,  $\mathbf{r} = \mathbf{v} = [0, \dots, 0] \in \mathbb{R}^d$ 
Input:  $\mathbf{X}, \mathbf{y}$ 
Output:  $\mathbb{S} \in \mathbb{R}^p$ ,  $0 < p \leq d$ 
1 for  $j \leftarrow 1$  to  $d$  do
2    $\bar{y} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$ 
3    $\bar{x}_j = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_{ij}$ 
4    $\theta_j = \frac{\sum_{i=1}^n (\mathbf{y}_i - \bar{y})(\mathbf{X}_{ij} - \bar{x}_j)}{\sum_{i=1}^n (\mathbf{X}_{ij} - \bar{x}_j)^2}$ 
   //  $r_j$  is the r-squared of a regression of  $\mathbf{y}$  onto  $\mathbf{X}_j$ .
5    $r_j = R_{\mathbf{y}|\mathbf{X}_j}^2 = 1 - \frac{\sum_{i=1}^n (\mathbf{y}_i - \theta_j \mathbf{X}_{ij})^2}{\sum_{i=1}^n (\mathbf{y}_i - \bar{y})^2}$ 
   //  $R_{\mathbf{X}_j|\mathbf{X}_{-j}}^2$  denotes the r-squared of a regression of  $\mathbf{X}_j$  onto the
   // remaining predictors, excluding the  $j'$ th one.
6    $\mathbf{v}_j = \frac{1}{1 - R_{\mathbf{X}_j|\mathbf{X}_{-j}}^2}$ 
7 end
8 Sort  $\mathbf{r}$  from smallest to largest and use its sort index to sort  $\mathbb{S}$ .
9 while  $\forall \mathbf{v} \geq \kappa$  do
10   for  $j \leftarrow 1$  to  $m$  do
11     if  $\mathbf{v}_j \geq \kappa$  then
12       remove  $\mathbb{S}_j$ 
13       recompute  $\mathbf{v}$ 
14     end
15   end
16 end

```

2.4 Kernel Ridge Regression

The Ridge regression model of Hoerl (1962) minimises the cost function

$$J_\theta = \frac{1}{2} \sum_{i=1}^n (y_i - \boldsymbol{\theta}^T \mathbf{x}_i)^2 + \frac{1}{2} \sum_{j=1}^d \lambda \boldsymbol{\theta}_j^T \boldsymbol{\theta}_j. \quad (1)$$

The regression parameters can be estimated sequentially in an online manner

$$\hat{\boldsymbol{\theta}}_t = \left[\left(\sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^T \right) + \lambda \mathbf{I}_d \right]^{-1} \sum_{i=1}^t \mathbf{x}_i y_i. \quad (2)$$

Ridge regression is amenable to being kernelised. Following Murphy (2012), let us define a kernel function to be a real-valued function of two arguments, $\mathcal{K}(\mathbf{x}, \mathbf{x}') \in \mathbb{R}$, for $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $\mathcal{X}$ being some abstract space. Typically the kernel function is symmetric and non-negative. A suitable kernel function for measuring similarity between vectors, is the radial basis function

$$\mathcal{K}(\mathbf{x}, \mathbf{x}') = \exp \left(-\frac{1}{2} (\mathbf{x} - \mathbf{x}')^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \mathbf{x}') \right), \quad (3)$$

The scalar variance version

$$\mathcal{K}(\mathbf{x}, \mathbf{x}') = \exp \left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\sigma^2} \right), \quad (4)$$

is seemingly more popular, most likely as the full covariance matrix can show greater prediction variance when used to generate features for a supervised learner. Let us define the kernel Gram matrix

$$\mathbf{K} = \begin{pmatrix} \mathcal{K}(\mathbf{x}_0, \mathbf{x}_0) & \dots & \mathcal{K}(\mathbf{x}_0, \mathbf{x}_n) \\ & \ddots & \\ \mathcal{K}(\mathbf{x}_n, \mathbf{x}_0) & \dots & \mathcal{K}(\mathbf{x}_n, \mathbf{x}_n) \end{pmatrix}.$$

Equations 1 and 2 can be thought of as a primal optimisation problem. We are then able to make use of the so-called kernel trick, which replaces all inner products of the form $\langle \mathbf{x}, \mathbf{x}' \rangle$ with kernel functions $\mathcal{K}(\mathbf{x}, \mathbf{x}')$. In the dual optimisation problem, we define the dual variables as

$$\boldsymbol{\alpha} = (\mathbf{K} + \lambda \mathbf{I}_n)^{-1} \mathbf{y},$$

which allows us to rewrite the primal variables as

$$\boldsymbol{\theta} = \mathbf{X}^T \boldsymbol{\alpha} = \sum_{i=1}^n \alpha_i \mathbf{x}_i.$$

The test time predictive mean is thus

$$\hat{y}_t = \hat{f}(\mathbf{x}_t) = \boldsymbol{\theta}^T \mathbf{x}_t = \sum_{i=1}^n \alpha_i \mathbf{x}_i^T \mathbf{x}_t = \sum_{i=1}^n \alpha_i \mathcal{K}(\mathbf{x}_i, \mathbf{x}_t).$$

As described by Vovk (2013), kernel Ridge regression was first coined by Cristianini and Shawe-Taylor (2000), and is a special case of support vector regression (Vapnik, 1998). Vovk goes on to show that kernel Ridge regression has certain performance guarantees that do not require any stochastic assumptions to be made. Specifically, he shows that for any sequence of time steps, the model satisfies

$$\sum_{t=0}^n \frac{(y_t - \hat{y}_t)^2}{1 + \frac{1}{\lambda} \mathcal{K}(\mathbf{x}_t, \mathbf{x}_t) - \frac{1}{\lambda} \mathbf{k}_t^T (\mathbf{K}_{t-1} + \lambda \mathbf{I}_n)^{-1} \mathbf{k}_t} = \min_{f \in \mathcal{F}} \left(\sum_{t=0}^n (y_t - \hat{y}_t)^2 + \lambda \|f\|_{\mathcal{F}}^2 \right),$$

where $\mathbf{k}_t$ is the vector with components

$$\mathbf{k}_t = \begin{pmatrix} \mathcal{K}(\mathbf{x}_0, \mathbf{x}_t) \\ \vdots \\ \mathcal{K}(\mathbf{x}_n, \mathbf{x}_t) \end{pmatrix}.$$

In addition, assuming $|y_t| \leq y_{max}$, and $\hat{y}_t$ is similarly clipped $|\hat{y}_t| \leq y_{max}$, then the sum of squared errors is bounded as:

$$\sum_{t=0}^n (y_t - \hat{y}_t)^2 \leq \min_{f \in \mathcal{F}} \left(\sum_{t=0}^n (y_t - \hat{y}_t)^2 + \lambda \|f\|_{\mathcal{F}}^2 \right) + 4y_{max}^2 \ln \det \left(\mathbf{I}_n + \frac{1}{\lambda} \mathbf{K}_t \right).$$

A practical difficulty with kernel Ridge regression, is that model fitting scales as $\mathcal{O}(n^3)$. Williams and Seeger (2001) show that an approximation to the eigen-decomposition of the Gram matrix can be computed by the Nyström method, which is used for the numerical solution of eigen-problems. This is achieved by carrying out an eigen-decomposition on a smaller system of size $m < n$, and then expanding the results back up to n dimensions. The computational complexity of a predictor using this approximation is $\mathcal{O}(m^2n)$. On experiments they conduct with a Gaussian process classifier, they demonstrate that the

Nyström approximation to the kernel Gram matrix allows a very significant speed-up of computation, without sacrificing accuracy. Pseudo-code for the algorithm is shown next.

Algorithm 2: the Nyström approximation to the kernel Gram matrix

Require: p, m , with $p \leq m < n$

Input: $\mathbf{X} \in \mathbb{R}^{n \times d}$

Output: $\mathbf{Z} \in \mathbb{R}^{m \times p}$, typically with $d < p$

- 1 Construct the kernel Gram matrix $\tilde{\mathbf{K}} \in \mathbb{R}^{m \times m}$ using random sampling.
 - 2 Compute a singular value decomposition $\tilde{\mathbf{K}} = \tilde{\mathbf{U}}\tilde{\mathbf{\Lambda}}\tilde{\mathbf{U}}^T$, where $\tilde{\mathbf{U}}$ is orthonormal, and $\tilde{\mathbf{\Lambda}}$ is a diagonal matrix of ranked eigenvalues, with $\tilde{\lambda}_0 \geq \tilde{\lambda}_1 \geq \dots \geq 0$.
 - 3 Take the first p columns of $\tilde{\mathbf{U}}$ and the first p eigenvalues of $\tilde{\mathbf{\Lambda}}$ to form $\tilde{\mathbf{M}} = \tilde{\mathbf{U}}_{:p}\tilde{\mathbf{\Lambda}}_{:p}$, with $\tilde{\mathbf{M}} \in \mathbb{R}^{m \times p}$.
 - 4 Reconstitute the kernelised predictor matrix $\mathbf{Z} = \tilde{\mathbf{K}}\tilde{\mathbf{M}}$, and pass this to the supervised learner of your choice.
-

3. The Radial Basis Function Network

The radial basis function network is a single layer network, where the hidden processing units play the role of $\mathbf{K}$, the kernel Gram matrix of kernel Ridge regression, see figure 1. However, the number of hidden processing units k in the radial basis function network, is usually $k \ll n$. The hidden processing units are commonly estimated via a clustering algorithm, rather than randomly selecting a subset of n training vectors as per the Nyström approximation to the kernel Gram matrix. Following Bishop (1995), the k^{th} output of the radial basis function network is defined as:

$$y_k(\mathbf{x}) = \sum_{j=1}^m \theta_{kj} \phi_j(\mathbf{x}) + \theta_{k0},$$

where θ_{kj} is the weight going from the j^{th} basis function (hidden processing unit) to the k^{th} output. In the case of a real-valued output, $k = 1$. The bias θ_{k0} can be absorbed into the summation by defining an extra basis function ϕ_0 with an activation set to 1. Nonlinearity is introduced into the network via the Gaussian basis function

$$\phi_j(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x} - \boldsymbol{\mu}_j\|_2^2}{2\sigma_j^2}\right),$$

where $\mathbf{x}$ is a d -dimensional input vector, $\boldsymbol{\mu}_j$ is the centre of the basis function $\phi_j(\cdot)$ and σ_j^2 is its width. Let Φ denote the set of hidden processing units with functions $\phi_j(\mathbf{x})$ that have been aggregated together. Bishop shows that the width σ_j^2 can be replaced by a covariance matrix $\boldsymbol{\Sigma}_j$, leading to basis functions of the form

$$\phi_j(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_j)^T \boldsymbol{\Sigma}_j^{-1}(\mathbf{x} - \boldsymbol{\mu}_j)\right). \quad (5)$$

Figure 2 demonstrates a comparison of a several classifiers on synthetic datasets. As Bishop (1995) discusses, a multi-layer perceptron separates classes using hidden units which form hyperplanes in the input space. Alternatively, separation of class distributions can be modelled by local radial basis kernel functions. The activations of the radial basis functions can be interpreted as the posterior probabilities of the presence of corresponding features in the input space, and the weights can be interpreted as the posterior probabilities of class membership, given the presence of the features.

Figure 1: architecture of the radial basis function network

Figure 2: comparison of a several scikit-learn classifiers on synthetic datasets

3.1 Literature Review

Girosi and Poggio (1990) show that radial basis function networks have a "best approximation" property, that is, in the set of approximating functions corresponding to all possible choices of parameters, there is one function which has minimum approximating error. Seminal papers on radial basis function networks include those of Broomhead and Lowe (1988) and Moody and Darken (1989). Moody and Darken's approach to network training, is extremely fast. It involves an unsupervised learning step, using a clustering algorithm

such as k-means (Lloyd, 1982) to learn the hidden processing unit means μ , and a supervised learning model such as Ridge regression which maps the hidden processing units Φ , to the output, y . They make use of a scalar variance term as in equation 4, rather than the full covariance as per equation 3. This method involves computing a "global first nearest-neighbor" heuristic, which uses a uniform average width $\sigma = \langle \Delta x_{\alpha\beta} \rangle$ for all hidden processing units, where $\Delta x_{\alpha\beta}$ is the Euclidean distance in input space between unit α and its nearest-neighbor β , and $\langle \rangle$ indicates a global average over such pairs. The fitting complexity for Moody and Darken's algorithm is $\mathcal{O}(knr)$ for the kmeans part, where k is the number of clusters and r is the number of fitting iterations, and $\mathcal{O}(n^2p + p^3)$ for the linear regression part.

Billings et al. (1996) consider a modified method of Moody and Darken's to network training. They consider a large input space in $\mathbf{X}$ and a large hidden network space in μ , relying on orthogonal least squares (Markovsky and Van Huffel, 2007) and forward stepwise selection to select the hidden processing units. Finally, they employ recursive least squares (Harvey, 1993) to map the hidden processing units to the response on an online basis. They make the link between the nonlinear autoregressive moving average model with exogenous inputs (Chen and Billings, 1989) and the radial basis function network, demonstrating an application to multiple input multiple output modelling (Bontempi, 2008) in simulated dynamic time series. A downside of their approach is that there is likely to be a lot of redundancy in the input space, and in practical real-time application, it may be wasteful to compute many predictors, only to potentially throw them away during network training.

Kanazawa (2020) applies an offline radial basis function network based on Moody and Darken's technique, to macroeconomic data. He finds that the estimated impulse responses from the model, suggest that the response of macroeconomic variables to a positive supply shock, is substantially time variant. He also finds that the model outperforms benchmarks based on the vector autoregression and threshold-var estimators, but only with longer term forecasts, 10 steps ahead or more. Overall, he finds that the model can uncover the structure of data generated from the nonlinear new Keynesian model, even from a small sample of simulated data. We can draw parallels between this outcome and the concept of few shot learning (Wang et al., 2017), where the learner can rapidly generalize to new tasks containing only a few samples of supervised information. He employs the renormalised radial basis function approach of Hastie et al. (2009)

$$h_c(x) = \frac{\phi_c(\mathbf{x})}{\sum_{j=1}^k \phi_j(\mathbf{x})},$$

which aims to fill any holes in regions of $\mathbb{R}^d$ where the kernels have no appreciable support.

Khosravi (2011) demonstrates an approach to radial basis function network training which is similar to backpropagation for neural networks (Rumelhart et al., 1986), although he sets some weights between the inputs and hidden layer, rather than the traditional approach, which has weights from the hidden layer to the outputs. He calls this his weighted rbfnet, and finds improved accuracy on the UCI letter classification dataset and the HODA digit recognition dataset.

The radial basis function network relates to the relevance vector machine of Tipping (1999). Originally, Tipping created this model as a sparse, Bayesian alternative to the support vector machine. The sparsity is induced by defining an automatic relevance determination Gaussian prior (MacKay 1994, Neal 1996) over the model weights. The model parameters are estimated via iterative reestimation of the individual weight priors α , through expectation maximisation (Dempster et al., 1977). Rasmussen and Quinonero-candela (2005) make the link between the relevance vector machine and the Gaussian process model, where the former's hyper-parameters are parameters of the latter's covariance function. Being a

local learning technique, Rasmussen and Quiñonero-candela highlight that whilst the relevance vector machine provides full predictive distributions for test cases, the predictive uncertainties have the unintuitive property that they get smaller the further you move away from the training cases. They propose to augment the relevance vector machine by an additional basis function centered at the test input, which adds extra flexibility at test time, thus improving generalisation performance. Gaussian process regression has an $\mathcal{O}(n^3)$ time complexity and an $\mathcal{O}(n^2)$ memory complexity, and is thus more computationally expensive to fit than the radial basis function network formulated by Moody and Darken. Furthermore, Moody and Darken’s radial basis function network is in effect making use of transfer learning: the knowledge of the clustering model is transferred to the supervised learner, and the intrinsic nature of the feature space is learnt and made available to the upstream model.

3.2 Clustering Algorithms

Assuming the Moody and Darken approach to radial basis function network formulation, there are several clustering algorithms that are suitable for use in deriving the hidden processing units. We discuss a few of these, which meet the following criteria:

1. The clustering algorithm can provide hidden processing unit means $\boldsymbol{\mu}_0, \dots, \boldsymbol{\mu}_k$.
2. Covariance matrices $\boldsymbol{\Sigma}_0, \dots, \boldsymbol{\Sigma}_k$ and their inverses (precision matrices) can be estimated and associated with each cluster mean.
3. The cluster means and covariances can be estimated online, sequentially.
4. The number of clusters, and in effect the network size of the radial basis function network, can be estimated quantitatively, without a more expensive estimation procedure such as cross-validation.

All the clustering algorithms discussed next, can have their means updated sequentially in a manner similar to equation 6. The precision matrices, which are required in equation 5, can be estimated in a manner similar to algorithm 4, or with exponential decay as in algorithm 3. All that is required is for the clustering algorithm or pseudo-clustering algorithm to make an assignment to the c' th cluster at test time.

3.2.1 KMEANS

In the classical kmeans algorithm, one selects k cluster centres a priori, and the training data $\mathbf{x} \in \mathbb{R}^d$ is assigned to the nearest cluster centre, with a goal of minimising the sum of squared distances. Usually these k centres are initialised at random. Each training vector is then assigned to the nearest center, and each center is recomputed as the mean of all points assigned to it. These two steps of assignment and mean calculation, are repeated until the process stabilises, or a maximum number of iterations is exceeded. The training error is

$$J_k = \sum_{j=1}^k \sum_{i=1}^n \delta_{ji} \|\mathbf{x}_i - \boldsymbol{\mu}_j\|_2^2.$$

Here δ_{ji} is 1 when the training exemplar $\mathbf{x}_i$ belongs to the processing unit $\boldsymbol{\mu}_j$, and 0 otherwise. A further advantage of this approach is that it is amenable to online learning. A partial or online update takes the form

$$\Delta \boldsymbol{\mu}_j = \eta(\mathbf{x}_i - \boldsymbol{\mu}_j), \tag{6}$$

where η is a learning rate in $[0, 1]$. Denote the optimal, minimal training error as J_k^* . Due to 'unlucky' random initialisation of the cluster means, the ratio $\frac{J_k}{J_k^*}$ is unbounded, even for fixed n and k . Arthur and Vassilvitskii (2007) demonstrate a way of initializing k-means by choosing random starting centers with very specific probabilities. They select a point i as a centre with probability proportional to the overall potential. Let $D(x)$ denote the shortest distance from a data point $\mathbf{x}$ to the closest centre already chosen. The next centre, denoted as $\boldsymbol{\mu}_c = \mathbf{x}'$, is then chosen with probability $\frac{D(\mathbf{x}')^2}{\sum_{i=1}^n D(\mathbf{x}_i)^2}$. They are able to upper bound the loss of their so-called kmeans++ algorithm as

$$\mathbb{E}[J_k] \leq 8(\ln k + 2)J_k^*.$$

Having estimated the cluster means $\boldsymbol{\mu}_j$, we are able to extract the cluster covariances

$$\boldsymbol{\Sigma}_j = \frac{1}{n_k - 1} \sum_{j=1}^k \sum_{i=1}^n \delta_{ji} (\mathbf{x}_i - \boldsymbol{\mu}_j)(\mathbf{x}_i - \boldsymbol{\mu}_j)^T, \quad (7)$$

where $n_k = \sum_{j=1}^k \sum_{i=1}^n \delta_{ji}$. How might we select the number of clusters k in a fast, quantitative manner? The silhouette score of Rousseeuw (1987) provides one solution. Denote as a_i the average dissimilarity of datum i to all other objects clustered in A , c_i the average dissimilarity of i to all other objects clustered in C , and $b_i = \min_{C \neq A} d_i$.

The silhouette score for the i 'th datum is

$$s_i = \begin{cases} 1 - a_i/b_i & \text{if } a_i < b_i \\ 0 & \text{if } a_i = b_i \\ b_i/a_i - 1 & \text{if } a_i > b_i \end{cases},$$

or equivalently

$$s_i = \frac{b_i - a_i}{\max(a_i, b_i)}.$$

Thus $-1 \leq s_i \leq 1$. Averaging the s_i over all samples allows us to select the number of clusters k such that

$$k^* = \max(\bar{s}_0, \dots, \bar{s}_J). \quad (8)$$

The approach has time complexity of $\sum_j^J \mathcal{O}(k_j n r)$, where k_j is the number of clusters for the j 'th clustering, n the number of training examples and r is the number of fitting iterations.

3.2.2 GAUSSIAN MIXTURE MODELS

Gaussian mixture models facilitate a probabilistic, parametric based approach to clustering, where the data generating process is assumed to be a mixture of multivariate Normal densities. Denote the probability density function of a k component mixture as

$$p(\mathbf{x}|\boldsymbol{\theta}) = \sum_{j=1}^k \pi_j p(\mathbf{x}|\boldsymbol{\theta}_j) = \sum_{j=1}^k \pi_j \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j),$$

where

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right], \quad (9)$$

and the mixing weights satisfy $0 \leq \pi_j \leq 1$, $\sum_{j=1}^k \pi_j = 1$. The maximum likelihood estimate

$$\boldsymbol{\theta}_{ML} = \arg \max_{\boldsymbol{\theta}} \ln p(\mathbf{x}|\boldsymbol{\theta}),$$

and the Bayesian maximum a posteriori criterion

$$\boldsymbol{\theta}_{MAP} = \arg \max_{\boldsymbol{\theta}} \ln p(\mathbf{x}|\boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}),$$

cannot be found analytically. The standard way of estimating $\boldsymbol{\theta}_{ML}$ or $\boldsymbol{\theta}_{MAP}$ is the expectation-maximisation algorithm (Dempster et al., 1977). This iterative procedure is based on the interpretation that $\mathbf{x}$ is incomplete data. The missing part for finite mixtures is the set of labels $\mathcal{Z} = \mathbf{z}_0, \dots, \mathbf{z}_n$ which accompany the training data $\mathbf{x}_0, \dots, \mathbf{x}_n$, indicating which component produced each training vector. Following Murphy (2012), let us define the complete data log likelihood to be

$$\ell_c(\boldsymbol{\theta}) = \sum_{i=1}^n \ln p(\mathbf{x}_i, \mathbf{z}_i|\boldsymbol{\theta}),$$

which cannot be computed, since $\mathbf{z}_i$ is unknown. Thus, let us define an auxiliary function

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}_{t-1}) = \mathbb{E}[\ell_c(\boldsymbol{\theta})|\mathbf{x}, \boldsymbol{\theta}_{t-1}],$$

where t is the current time step. The expectation is taken with respect to the old parameters $\boldsymbol{\theta}_{t-1}$ and the observed data $\mathbf{x}$. Denote as $r_{ic} = p(z_i = c|\mathbf{x}_i, \boldsymbol{\theta}_{t-1})$, the responsibility that cluster c takes for datum i . The expectation step has the following form

$$r_{ic} = \frac{\pi_c p(\mathbf{x}_i|\boldsymbol{\theta}_{c,t-1})}{\sum_{j=1}^k \pi_j p(\mathbf{x}_i|\boldsymbol{\theta}_{j,t-1})}.$$

The maximisation step optimises the auxiliary function $\mathcal{Q}$ with respect to $\boldsymbol{\theta}$

$$\boldsymbol{\theta}_t = \arg \max_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}_{t-1}).$$

The c 'th mixing weight is estimated as

$$\pi_c = \frac{1}{n} \sum_{i=1}^n r_{ic} = \frac{r_c}{n}.$$

The parameter set $\boldsymbol{\theta}_c = \{\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c\}$ is then

$$\begin{aligned} \boldsymbol{\mu}_c &= \frac{\sum_{i=1}^n r_{ic} \mathbf{x}_i}{r_c} \\ \boldsymbol{\Sigma}_c &= \frac{\sum_{i=1}^n r_{ic} (\mathbf{x}_i - \boldsymbol{\mu}_c)(\mathbf{x}_i - \boldsymbol{\mu}_c)^T}{r_c}. \end{aligned}$$

As discussed by Figueiredo and Jain (2002), expectation-maximisation is highly dependent on initialisation. They highlight several strategies to ameliorate this problem, such as multiple random starts, with final selection based on the highest maximum likelihood of the mixture, or kmeans based initialisation. However, in mixture models, the distinction between model-class selection and model estimation is unclear. For example, a 3 component mixture in which one of the mixing probabilities is zero, is indistinguishable for a 2 component mixture. They propose an unsupervised algorithm for learning a finite

mixture model from multivariate data. Their approach is based on the philosophy of minimum message length encoding (Wallace and Dowe, 1999), where one aims to build a short code that facilitates a good data generation model. Their algorithm is capable of selecting the number of components and unlike the standard expectation-maximization algorithm, does not require careful initialization. The proposed method also avoids another drawback of expectation-maximization for mixture fitting: the possibility of convergence toward a singular estimate at the boundary of the parameter space. Denote the optimal mixture parameter set

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \ell_{FJ}(\boldsymbol{\theta}, \mathbf{x}),$$

where

$$\ell_{FJ}(\boldsymbol{\theta}, \mathbf{x}) = \frac{n}{2} \sum_{j=1}^k \ln \left(\frac{n\pi_k}{12} \right) + \frac{k}{2} \ln \left(\frac{n}{12} \right) + \frac{k(n+1)}{2} - \ln p(\mathbf{x}|\boldsymbol{\theta}).$$

This leads to a modified maximisation step

$$\pi_c = \frac{\max \left\{ 0, \left(\sum_{i=1}^n r_{ic} \right) - \frac{n}{2} \right\}}{\sum_{j=1}^k \max \left\{ 0, \left(\sum_{i=1}^n r_{ij} \right) - \frac{n}{2} \right\}}$$

for $c = 1, 2, \dots, k$.

The maximisation step is identical to expectation-maximisation, except that the c 'th parameter set $\boldsymbol{\theta}_c$ is only estimated when $\pi_c > 0$, and $\boldsymbol{\theta}_c$ is discarded from $\boldsymbol{\theta}^*$ when $\pi_c = 0$. A distinctive feature of the modified maximisation step is that it leads to component annihilation. This prevents the algorithm from approaching the boundary of the parameter space. In other words, if one of the mixtures is not supported by the data, it is annihilated.

3.2.3 DISCRIMINANT ANALYSIS AS A PSEUDO-CLUSTERING ALGORITHM

When we consider the four selection criteria stated at the start of section 3.2, quadratic discriminant analysis, where we must estimate class-conditional means and covariances, modelled by multivariate normal distributions, could be considered as a pseudo-clustering algorithm. Instead of an unsupervised learning step to infer labels, we can create labels from the regression responses, and derive the class-conditional parameter sets $\boldsymbol{\theta}_c = \{\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c\}$. For financial time series, an obvious set of classes is $\{-1, 0, 1\}$, corresponding to the signed mid to mid returns. There is a sound, scientific rationale for deriving class-conditional covariates for such time series. It is well known amongst financial practitioners that downside volatility of returns behaves differently to up-side volatility, particularly for equities. In the plot and tabular summary shown next, we display annualised volatility (standard deviations) computed from 21 day sliding windows of daily mid to mid returns for the S&P 500 equity index. The data are extracted from Refinitiv. The volatilities are separated out by negative and positive daily returns. We see that since the 1987 stock market crash, so-called 'Black Monday', annualised volatility is around 0.8 of a percent higher for downside returns.

regime	count	mean	std	min	q25	q50	q75	max	ann vol
-1	11219	-0.78%	0.92%	-20.47%	-1.00%	-0.51%	-0.22%	0.00%	14.60%
1	12696	0.74%	0.86%	0.00%	0.23%	0.50%	0.95%	16.61%	13.72%

Table 1: .SPX annualised volatility by regime

Figure 3: .SPX annualised volatility by regime

3.3 Regularisation Of Covariance Matrices And Their Inverses

A central issue for the algorithms discussed in sections 3.2.1, 3.2.2 and 3.2.3, is the estimation of covariance matrices, that are required by the radial basis function network. As highlighted by Friedman (1989), equation 7 produces biased estimates of the eigenvalues; the largest ones are biased high and the smallest ones are biased toward values that are too low. He goes on to say that the net effect of this biasing phenomenon on discriminant analysis is to (sometimes dramatically) exaggerate the importance associated with the low-variance subspace spanned by the eigenvectors corresponding to the smallest sample eigenvalues. Aside from this, we need a way to reduce or mitigate completely, the numerical issues that appear when estimating covariance matrices where $n < d$, the number of predictors is larger than the number of observations. Friedman's procedure to regularise covariance matrices involves two steps. Let us define the pooled covariance estimate as Σ , which is estimated from all the training data. In the first step, we shrink the class conditional covariance toward the pooled estimate

$$\Sigma_c(\lambda) = (1 - \lambda)\Sigma_c + \lambda\Sigma,$$

where $0 \leq \lambda \leq 1$. In the second step, we shrink $\Sigma_c(\lambda)$ toward a multiple of the identity matrix

$$\Sigma_c(\lambda, \gamma) = (1 - \gamma)\Sigma_c(\lambda) + \frac{\gamma}{d}\text{trace}[\Sigma_c(\lambda)]\mathbf{I},$$

with $0 \leq \gamma \leq 1$. A standard way to estimate λ and γ is through cross-validation. Another area in which stability can be gained, is to estimate the determinant of the covariance

matrix of equation 9, via a spectral decomposition

$$\Sigma_c = \sum_{j=1}^k e_{jc} \mathbf{v}_{jc} \mathbf{v}_{jc}^T,$$

where e_{jc} is the j' th eigenvalue of Σ_c in decreasing value, and $\mathbf{v}_{jc}$ is the corresponding eigenvector. The stabilised determinant of the covariance matrix for the multivariate normal is then

$$|\Sigma_c| = \sum_{j=1}^k \max(0, e_{jc}).$$

3.4 Exponentially Weighted Recursive Least Squares

Once we have selected the hidden processing unit centres and covariances using any of the algorithms in section 3.2, we must then map this new predictor space to the regression target. A suitable supervised learning algorithm, is the recursive least squares estimator, which is a special case of the Kalman filter (Harvey, 1993). In particular, we are interested in the exponentially weighted formulation, see Liu et al. (2010), which facilitates adaptation sequentially in time. The algorithm we show below, includes a variance stabilisation update, which ameliorates the build-up of large values along the diagonal of the precision matrix $\mathbf{P}$, which may occur if the response y has low variance. See for example Gunnarsson (1996) for a further discussion on the regularisation of recursive least squares. Similar regularisation approaches suitable for online learning and nonstationary data are studied by Moroshko et al. (2015).

Algorithm 3: exponentially weighted recursive least squares

Require: λ // the Ridge penalty
1 $0 \ll \tau < 1$ // an exponential forgetting factor
Initialise: $\boldsymbol{\theta} = \mathbf{0}$, $\mathbf{P} = \mathbf{I}_d / \lambda$
Input: $\mathbf{x}_t \in \mathbb{R}^{d \times 1}$, y_t
Output: $\hat{y}_t$
2 $r = 1 + \mathbf{x}_t^T \mathbf{P}_{t-1} \mathbf{x}_t / \tau$
3 $\mathbf{k} = \mathbf{P}_{t-1} \mathbf{x}_t / (r\tau)$
4 $\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \mathbf{k}(y_t - \boldsymbol{\theta}_{t-1}^T \mathbf{x}_t)$
5 $\mathbf{P}_t = \mathbf{P}_{t-1} / \tau - \mathbf{k} \mathbf{k}^T r$
6 $\mathbf{P}_t = \mathbf{P}_t \tau$ // variance stabilisation
7 $\hat{y}_t = \boldsymbol{\theta}_t^T \mathbf{x}_t$

4. Ensemble Learning

We have shown that feature selection can be achieved through the meta-algorithm 1. In addition, the hidden processing units of the radial basis function network, with centres and covariances chosen by the clustering algorithms of section 3.2, can be thought of as a form of model selection. It is possible to add another layer of model selection, by combining the forecasts of the individual models. Here we explore two forms of such ensemble learning: weighted average forecasters and so-called follow the best expert ensembles. We begin first with some background material in the theory of model prediction with expert advice. Following Cesa-Bianchi and Lugosi (2006), let us define a sequential decision maker whose goal is to predict an unknown sequence $y_1, y_2, \dots$ of an outcome space $\mathcal{Y}$. The ensemble

forecaster's predictions $\hat{p}_1, \hat{p}_2, \dots$ belong to a decision space $\mathcal{D}$. After each forecast round, the predictive performance of each forecaster is compared to a set of reference experts, which we denote as $\hat{\mathbf{y}}_t = \hat{y}_{t,1}, \dots, \hat{y}_{t,m}$. Let us define a non-negative loss function $\ell : \mathcal{D} \times \mathcal{Y} \rightarrow \mathbb{R}$. The forecaster's goal is to minimise the cumulative regret with respect to each expert

$$R_{n,m} = \sum_{t=1}^n [\ell(y_t, \hat{p}_t) - \ell(\hat{y}_{t,m}, \hat{p}_t)] = \hat{L}_n - L_{n,m}.$$

The weighted average forecaster predicts at time t according to

$$\hat{p}_t = \frac{\sum_{i=1}^m w_{t-1,i} \hat{y}_{t,i}}{\sum_{j=1}^m w_{t-1,j}}.$$

Assuming ℓ is convex in its first argument, and that it takes values in $[0, 1]$, then after n prediction rounds, for any sequence $y_1, y_2, \dots, y_m \in \mathcal{Y}$, the regret of the weighted average forecaster satisfies

$$\hat{L}_n - \min_{j=1, \dots, m} L_{n,m} \leq \sqrt{nm}.$$

4.1 Precision Weighted Ensemble

A close relative to the weighted average forecaster, is the precision weighted ensemble. Let us construct a next step ahead forecaster, which is weighted by the inverse of the squared errors made by an ensemble of experts. Denote as before the expert forecasts as $\hat{\mathbf{y}}_t$, where t rounds have been observed. We can use the Sherman-Morrison formula (Sherman and Morrison, 1950), as shown by Duda et al. (2001), to derive a precision matrix sequentially. In our case, we are interested in a precision matrix of forecast errors. A complete algorithm is shown next.

Algorithm 4: precision weighted ensemble

Require: λ , a regularisation penalty, with $\lambda \geq 0$
Initialise: $\mathbf{P}_0 = \mathbf{I}/\lambda$
Input: $y_t, \hat{\mathbf{y}}_t$
Output: $\hat{p}_t$

- 1 $\mathbf{e}_t = y_t - \hat{\mathbf{y}}_{t-1}$
- 2 $\bar{\mathbf{e}}_t = \frac{t-1}{t} \bar{\mathbf{e}}_t + \frac{1}{t} \mathbf{e}_t$
- 3 $\mathbf{A} = \mathbf{P}_{t-1}(\mathbf{e}_t - \bar{\mathbf{e}}_t)(\mathbf{e}_t - \bar{\mathbf{e}}_t)^T \mathbf{P}_{t-1}$
- 4 $\mathbf{B} = (\mathbf{e}_t - \bar{\mathbf{e}}_t)^T \mathbf{P}_{t-1}(\mathbf{e}_t - \bar{\mathbf{e}}_t)$
- 5 $c = \frac{t^2}{t+1}$
- 6 $\mathbf{P}_t = \frac{t+1}{t} \mathbf{P}_{t-1} - \frac{\mathbf{A}}{\mathbf{B}+c}$
- 7 $\mathbf{w}_t = \text{diag}(\mathbf{P}_t) / \text{trace}(\mathbf{P}_t)$
- 8 $\hat{p}_t = \mathbf{w}_t^T \hat{\mathbf{y}}_t$

Financial time series are generally non-stationary, as indicated by the jump-diffusion process hypothesis. We can swap out the simple moving average update of line 2 in algorithm 4, for an exponential update

$$\bar{\mathbf{e}}_t = \bar{\mathbf{e}}_{t-1} - \eta(\bar{\mathbf{e}}_{t-1} - \mathbf{e}_t),$$

where $0 < \eta \leq 1$. In addition, the precision matrix could be updated in a similar manner to exponentially weighted recursive least-squares, algorithm 3.

4.2 Follow The Best Expert

A simple forecasting strategy that is available to an ensemble, is one in which the expert that minimises the cumulative loss over the past $t - 1$ rounds, is followed at time t . As discussed by Cesa-Bianchi and Lugosi (2006), consider a class of $\mathcal{M}$ experts, and define the forecaster that predicts the same as the expert that minimises the cumulative historical loss

$$\hat{p}_t = \hat{y}_{t,m} \text{ if } m = \arg \min_{m' \in \mathcal{M}} \sum_{i=1}^{t-1} \ell(\hat{y}_{i,m'}, y_i).$$

Consider as well, the forecaster with forward looking bias

$$p_t^* = \hat{y}_{t,m} \text{ if } m = \arg \min_{m' \in \mathcal{M}} \sum_{i=1}^t \ell(\hat{y}_{i,m'}, y_i).$$

In the case of a square loss

$$\ell(y, p) = \|y - p\|^2,$$

we obtain the simple solution

$$\hat{p}_t = \frac{1}{t-1} \sum_{i=1}^{t-1} y_i,$$

and

$$p_t^* = \frac{1}{t} \sum_{i=1}^t y_i.$$

Then, for any $y \in \mathcal{D}$,

$$\ell(y, \hat{p}_t) - \ell(y, p_t^*) = \|y - \hat{p}_t\|^2 - \|y - p_t^*\|^2 \leq 4 \|\hat{p}_t - p_t^*\|.$$

For the accumulated loss after n rounds, we have the upper bound regret

$$\hat{L}_n - \inf_{m \in \mathcal{M}} L_{n,m} \leq \sum_{t=1}^n \frac{8}{t} \leq 8(1 + \ln n).$$

5. The Research Experiment

We consider the goal of multi-step forecasting with financial time series. Multi-step forecasting has practical use within electronic trading. The market maker isn't certain of when his risk might be hedged as he is uncertain as to when the liquidity takers will trade. The speculator incurs uncertainty too. She is likely to trade when her predictive signal indicates a larger potential profit than the cost of executing a trade. Common to both is the variable holding times of risk. With this in mind, multi-step forecasting provides a benefit when the optimal forecast horizon is unknown a priori. We begin with a description of the data and the models that we use in our research experiment, and follow this up with a description of the experiment's design.

5.1 Data

Refinitiv, formerly Thomson Reuters, is a global provider of financial market data and infrastructure. We extract minutely sampled data for various asset classes, including currency pairs, equities, interest rates, credit, metals, commodities, energy and crypto. These

are predominately cash instruments, but also include futures. The sampled prices are usually the last traded price in that interval. In some cases they are last quoted limit order book prices observed in that time interval. So as to ensure the largest amount of product coverage, we include instruments for which their exchange trading session is open between 8am and 4pm GMT. Thus, for example, we have to exclude some of the available Asian and American exchange traded instruments. Refinitiv restricts us to around 40000 historical quotes per Refinitiv information code (ric), which we invariably must trim back to meet the opening hours we are interested in. All said, our experiment includes roughly one month of minutely sampled data. Appendix A provides details of the rics that we are able to use.

5.2 Offline Learning Models

The offline learning models are:

- random walk model - this is the baseline model, described in section 2.1.
- ar(1) - the autoregressive order 1 model described in section 2.2.
- ridge - the Ridge regression model described in section 2.4.
- kernel ridge - the kernel Ridge regression model described in section 2.4. Here we use the Nyström approximation to the kernel Gram matrix, algorithm 2.

5.3 Online Learning Models

Note that with all the radial basis function networks shown next, the hidden processing units are mapped to the response via exponentially weighted recursive least-squares, algorithm 3. All covariance matrices are regularised as per Friedman (1989)’s paper, as discussed in section 3.3. The online learning models are:

- ewrls - an exponentially weighted recursive least-squares model.
- rbfnet km - the radial basis function network, with hidden processing units formed by the kmeans++ algorithm as described in section 3.2.1. Note that the number of means and associated precision matrices are chosen by the silhouette score of equation 8.
- rbfnet gmm - a radial basis function network formed by Gaussian mixture models as outlined in section 3.2.2. Specifically, we select the number of mixtures, and therefore the number of hidden processing units, as per Figueiredo and Jain (2002)’s algorithm.
- rbfnet rda - a radial basis function network formed of regularised quadratic discriminant analysis components. The labelled classes are inferred as per section 3.2.3.
- pwe - all online learning models form an ensemble of experts which are weighted by algorithm 4.
- fte - all online learning models form an ensemble of experts, in which all of the weight is given to the best expert that is followed, as described in section 4.2.

5.4 Experiment Design

These are the steps required in order to replicate and conduct the experiment.

1. The first step is to select a set of candidates from Refinitiv, from which we can extract historical, minutely sampled data, and whose quotes are active during GMT 8am to 4pm.

2. The second step is to set various hyper-parameters for the experiment, which are fixed in order to minimise computational cost. These include a maximum multi-step forecast horizon of 15 minutes, an exponential decay parameter $\tau = 0.999$ for the exponentially weighted recursive least-squares algorithm 3, and a maximum variance inflation factor $\kappa = 5$ for feature selection algorithm 1. We also fix the covariance shrinkage penalty $\lambda = 0.001$ and eigenvalue debiasing $\gamma = 0.001$ for the regularisation of covariance matrices as discussed in section 3.3.
3. Half the data is set aside for offline training, and the other half for testing. The online learning models use all the available information up to time $t - 1$ in the test set, when making forecasts at time t .
4. In the offline learning phase, initial feature selection is performed using the forward step-wise variance inflation factor minimisation algorithm 1.
5. For all the models that require Ridge penalties, we split the training set into randomly sampled training and validation subsets. In the training subset, we fit the models with varying Ridge penalties, using the analytical least-squares solution, equation 2. We then select the models whose Ridge penalties minimise the validation subset generalised cross-validation error

$$h = \frac{1}{d} \text{diag} \left(\mathbf{X} [\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I}_d]^{-1} \mathbf{X}^T \right)$$

$$gcv = \frac{1}{n-1} \sum_{i=1}^{n-1} \left(\frac{y_{i+1} - \hat{y}_i}{1-h} \right)^2.$$

6. During offline training, the structure of the various radial basis function networks is determined by the underlying clustering algorithms which are discussed in section 3.2.
7. Once all the online learning models are fitted, they are aggregated and made available to the two ensemble learners of section 4.
8. Finally, individual model performance is evaluated in the test set using a normalised mean-square prediction error. Let us define the model's mean-square prediction error for the h' th forecast horizon as

$$mspe_{model,h} = \frac{1}{t-h} \sum_{i=1}^{t-h} (y_{i+h} - \hat{y}_i)^2.$$

The normalised mean square prediction error for the h' th forecast horizon, where the normalisation occurs with respect to the random walk baseline model of section 2.1, is thus

$$nmspe_{model,h} = \frac{\frac{1}{t-h} \sum_{i=1}^{t-h} (y_{i+h} - \hat{y}_i)^2}{\frac{1}{t-h} \sum_{i=1}^{t-h} (y_{i+h})^2}.$$

One interpretation of the normalised mean-square prediction error, is the percentage improvement in accuracy over the baseline model in predicting the response.

6. Results

The tables that follow, show normalised mean square prediction error by model and forecast horizon. If this measure is below 1, then we outperform the random walk baseline. If it is above 1, then we underperform the baseline. We also perform a Wald test for the null hypothesis that the normalised mean square prediction error is no different from 1, which is tested at the 5% critical value. We see that several patterns emerge. Firstly, when considering the model averaged normalised mean-square prediction errors, none of

the offline learning models perform better than the random walk baseline. Secondly, feature selection has helped these offline models perform better than the $\text{ar}(1)$ offline model. We cannot however see a benefit of nonlinear modelling in the form of kernel Ridge regression over regular Ridge regression here. Thirdly, all online learning models outperform the baseline. The best performing models are the online radial basis function networks formed of $\text{kmeans}++$ and Gaussian mixture model clusters. The online radial basis function network formed of class-conditional means and covariances as per the regularised quadratic discriminant analysis algorithm outlined in section 3.2.3, performs less well than the online exponentially weighted recursive least-squares model. It is likely that too much learning capacity is taken out of this radial basis function network, as there are now just three hidden processing units corresponding to the signed mid-to-mid returns class labels $-1, 0, 1$. This model still performs better on average than the random walk baseline, although we cannot reject the hypothesis that the observed normalised mean square prediction error is no different from 1, according to the Wald test.

Both online ensembles outperform the random walk baseline. The precision weighted ensemble has slightly better results than the follow the expert ensemble. It is likely that the forecasts of the individual online models are quite highly correlated. This results in the ensembles performing slightly worse than the best experts. We suggest that the ensembles would have every chance of performing best if less-correlated, highly predictive experts were supplied to them. In many respects, this is the 'holy grail' of the prediction of sequences in financial time series, as experts are in fact often correlated. In table 3, we see the normalised mean square prediction error summarised by forecast horizon and model. There isn't a noticeable degradation or improvement over the random walk baseline, given a change in forecast horizon. In other words, the results across the various forecast horizons are stable. Finally, table 4 shows the percentage of experiments for which the individual models outperformed the random walk baseline. All online learning models outperform the baseline in over 70% of cases. The online radial basis function networks formed of $\text{kmeans}++$ and Gaussian mixture model clusters, are tied with exponentially weighted recursive least-squares at 76.8%. The precision weighted ensemble has the highest score, outperforming the baseline in 78% of all cases.

	ar(1)	kernel ridge	ridge	fte	pwe	rbfnet gmm	rbfnet km	rbfnet rda	ewrls
n_targets	82	82	82	82	82	82	82	82	82
count	1230	1230	1230	1230	1230	1230	1230	1230	1230
mean	2.055	1.942	1.850	0.897	0.865	0.834	0.833	0.986	0.870
std	2.220	1.659	1.424	0.570	0.537	0.459	0.461	0.758	0.560
min	0.392	0.253	0.264	0.197	0.222	0.237	0.232	0.201	0.221
25%	0.862	0.914	0.985	0.571	0.568	0.535	0.527	0.618	0.518
50%	1.475	1.543	1.415	0.763	0.770	0.747	0.747	0.819	0.731
75%	2.130	2.343	2.275	1.043	0.951	0.937	0.937	1.108	0.981
max	13.054	12.079	7.256	3.054	3.313	2.590	2.323	5.882	3.039
se	0.063	0.047	0.041	0.016	0.015	0.013	0.013	0.022	0.016
t-value	16.665	19.922	20.931	-6.351	-8.781	-12.656	-12.686	-0.665	-8.160
crit-value	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.253	0.000
reject H0	1	1	1	1	1	1	1	0	1

Table 2: summarised experiment results by model

h	ar(1)	kernel ridge	ridge	fte	pwe	rbfnet gmm	rbfnet km	rbfnet rda	ewrls
1	2.057	1.944	1.852	0.896	0.864	0.833	0.832	0.985	0.869
2	2.057	1.944	1.852	0.896	0.865	0.834	0.832	0.985	0.869
3	2.056	1.943	1.851	0.896	0.865	0.834	0.833	0.985	0.869
4	2.056	1.943	1.851	0.896	0.865	0.834	0.833	0.985	0.869
5	2.056	1.943	1.851	0.896	0.865	0.834	0.833	0.985	0.869
6	2.055	1.943	1.851	0.896	0.865	0.834	0.833	0.985	0.869
7	2.055	1.943	1.850	0.897	0.865	0.834	0.833	0.986	0.870
8	2.055	1.942	1.850	0.897	0.865	0.835	0.833	0.986	0.870
9	2.055	1.942	1.850	0.897	0.866	0.835	0.833	0.986	0.870
10	2.054	1.942	1.850	0.897	0.866	0.835	0.834	0.986	0.870
11	2.054	1.942	1.849	0.897	0.866	0.835	0.834	0.986	0.870
12	2.054	1.941	1.849	0.897	0.866	0.835	0.834	0.986	0.870
13	2.054	1.941	1.849	0.897	0.866	0.835	0.834	0.986	0.870
14	2.053	1.941	1.849	0.897	0.866	0.835	0.834	0.986	0.870
15	2.053	1.941	1.848	0.898	0.866	0.836	0.834	0.987	0.871

Table 3: summarised experiment results by horizon and model

model	outperformance
ridge	28.0%
ar(1)	29.3%
kernel ridge	29.3%
rbfnet rda	70.7%
fte	74.4%
rbfnet gmm	76.8%
rbfnet km	76.8%
ewrls	76.8%
pwe	78.0%

Table 4: percentage of cases where the baseline is outperformed

7. Discussion

Whilst the radial basis function networks formed of kmeans++ and Gaussian mixture models use more involved procedures, with the first clustering algorithm comparing multiple network structures and selecting the one which maximises a silhouette score, and the second algorithm starting with a large mixture and pruning it back until the change in modified log-likelihood function improves very little, these algorithms provide a quantitative basis for radial basis function network hidden processing unit selection. The meta-algorithms are still less expensive from a time complexity perspective to cross-validation, or even generalised cross-validation.

We settled on using exponentially weighted recursive least-squares to map the hidden processing units of the radial basis function network to the outputs. Although not shown in the final results, we did also experiment with lasso regression (Tibshirani, 1996). The specific algorithm employed was cyclic coordinate descent (Hastie et al., 2015). We found that lasso produced very little or no weight shrinkage of the weights going from the hidden

processing units to the outputs, indicating that each node in the network contains useful information in forecasting the response.

In terms of future work, one could consider using other unsupervised learning approaches to learn the hidden processing units of the radial basis function network. One such example is stacked denoising autoencoders (Vincent et al., 2010). Another option is variational Bayes expectation maximisation, which Murphy (2012) shows is sparsity promoting, typically requiring fewer clusters than the predefined set required for a Gaussian mixture model. Kingma and Welling (2014) combine autoencoders with variational Bayes.

More generally, we believe that continual learning with nonstationary time series is a relevant area of research. Although transfer learning is primarily concerned with transferring knowledge from a source domain to a target domain, and therefore may be used in either an offline or online setting, an increasing number of papers focus on online transfer learning. See for example Zhao et al. (2014), Salvalaio and de Oliveira Ramos (2019) and Wang et al. (2020). Our paper has contributed to the research of continual learning in financial time series. This is important as our experiment shows that continual learning provides a benefit with multi-step forecasting, above and beyond sequential learning. If we compare the local learning of the radial basis function network with the global learning technique of the feed-forward neural network, the latter suffers from catastrophic forgetting. Kirkpatrick et al. (2017) and Sukhov et al. (2020) look at ways of improving this issue, specifically at training networks that can maintain expertise on tasks that they have not experienced for a long time. The radial basis function network that we formulate, is naturally designed to measure the similarity between test samples and continuously updated prototypes that capture the characteristics of the feature space. As such, the model is somewhat robust in mitigating catastrophic forgetting.

8. Conclusion

Financial time series exhibit the attributes of autocorrelation, nonstationarity and nonlinearity. The experiment we conduct on minutely sampled datasets, demonstrate the added value of feature selection, nonlinear modelling and online learning, in the pursuit of providing multi-step forecasts of financial time series. We have experimented with offline learning models including an autoregressive(1), Ridge regression and Kernel Ridge regression. These models have been baselined against a random walk model, which they are unable to beat in multi-step forecasting of log returns. In contrast, the online learning models, which either use sequential updating alone in the form of an exponentially weighted recursive least-squares model, or a combination of sequential updating and transfer learning, in the form of radial basis function networks, are able to outperform the baseline. Finally, we experiment with a combination of experts that are tasked with the prediction of sequences, in the form of precision weighted and follow the best expert ensembles, which are also able to perform better than the baseline.

Acknowledgements

We would like to thank Stephen Roberts, RAEng / Man Group Professor of Machine Learning at Oxford University, and Professor Mark Herbster, programme director for the MSc in Machine Learning at University College London, for their thoughts on online learning. In addition, we would like to thank Professor Alessio Sancetta, Chair of Economics at Royal Holloway, for his suggestions that shaped the experiment toward an outcome that is more meaningful for the econometrics community.

Appendix A. Full Results

The full experiment results are shown in this section. Specifically we show the normalised mean-square prediction errors per target Refinitiv information code and model. There are 8,348 training observations and 8,347 test observations. The column $\mathbf{p}$ indicates the number of features chosen by the feature selection meta-algorithm shown in section 2.3. The $\text{ar}(1)$, kernel Ridge regression and Ridge regression models, are all fitted once, offline. The remaining models are fitting sequentially in the test set, where we make use of information up to time $t - 1$, when forecasting at time t . All radial basis function networks make use of algorithm 3 to map the hidden processing units to the targets.

ric	description	asset class	p	ar(1)	kernel ridge	ridge	fte	pwe	rbfnet gmm	rbfnet km	rbfnet rda	ewris
BTC=	Bitcoin/US Dollar	cash crypto	14	13.037	2.638	5.658	2.917	3.311	2.227	2.256	5.879	3.037
ETH=	Ethereum/US Dollar	cash crypto	11	2.03	2.739	2.654	0.752	0.708	0.77	0.743	1.016	0.751
LTC=	Litecoin/US Dollar	cash crypto	24	1.438	0.915	1.345	0.58	0.552	0.53	0.544	0.785	0.546
XRP=	Ripple/US Dollar	cash crypto	16	0.693	0.887	1.136	0.598	0.591	0.473	0.506	0.853	0.596
AUD=	Australian Dollar-US Dollar	cash FX	30	3.101	1.866	1.041	0.764	0.782	0.786	0.796	0.716	0.86
AUDCHF=	Australian Dollar-Swiss Franc	cash FX	41	1.216	0.885	1.583	0.419	0.403	0.412	0.409	0.455	0.387
AUDJPY=	Australian Dollar-Japanese Yen	cash FX	41	0.653	1.336	0.754	0.708	0.548	0.487	0.485	0.554	0.71
CAD=	US Dollar-Canadian Dollar	cash FX	24	8.023	2.889	0.741	1.243	0.951	0.851	0.759	1.051	1.303
CADCHF=	Canadian Dollar-Swiss Franc	cash FX	23	4.13	2.569	4.105	0.926	0.822	0.791	0.777	0.897	0.933
CADJPY=	Canadian Dollar-Japanese Yen	cash FX	20	1.212	2.039	0.722	0.301	0.318	0.353	0.329	0.283	0.321
CHF=	US Dollar-Swiss Franc	cash FX	19	2.608	0.956	3.32	1.091	0.859	0.711	0.763	1.11	0.903
CHFJPY=	Swiss Franc-Japanese Yen	cash FX	17	0.428	0.253	0.457	0.198	0.23	0.265	0.264	0.202	0.252
EUR=	Euro-US Dollar	cash FX	18	3.787	3.5	7.219	2.033	2.101	2.155	2.13	2.031	2.101
EURAUD=	Euro-Australian Dollar	cash FX	33	6.385	4.273	4.844	1.575	1.572	1.53	1.605	1.543	1.635
EURCAD=	Euro-Canadian Dollar	cash FX	18	0.713	1.837	1.917	0.852	1.07	1.294	1.197	0.847	0.985
EURCHF=	Euro-Swiss Franc	cash FX	28	0.866	0.902	1.07	1.303	1.134	1.081	1.072	1.303	1.091
EURGBP=	Euro-British Pound	cash FX	19	2.299	2.589	2.527	0.785	0.762	0.765	0.767	0.78	0.77
EURJPY=	Euro-Japanese Yen	cash FX	11	1.232	0.412	1.098	0.364	0.451	0.561	0.398	0.486	0.449
GBP=	British Pound-US Dollar	cash FX	15	7.057	4.842	2.415	2.775	2.456	2.048	2.068	3.274	2.563
GBPAUD=	British Pound-Australian Dollar	cash FX	31	1.462	4.239	2.187	0.794	0.86	0.711	0.828	1.011	0.964
GBPCAD=	British Pound-Canadian Dollar	cash FX	20	3.693	1.226	4.191	1.248	1.285	1.385	1.447	1.121	1.247
GBPCHF=	British Pound-Swiss Franc	cash FX	16	2.13	1.313	1.631	1.266	1.166	1.083	1.077	1.464	1.069
GBPJPY=	British Pound-Japanese Yen	cash FX	13	0.689	0.606	1.057	0.352	0.375	0.364	0.36	0.426	0.408
GBPNZD=	British Pound-New Zealand Dollar	cash FX	31	0.751	1.898	1.495	0.961	0.811	0.672	0.739	0.938	0.932
HKD=	US Dollar-Hong Kong Dollar	cash FX	11	1.059	0.478	1.346	0.372	0.331	0.526	0.313	0.34	0.295
INR=	US Dollar-Indonesia Rupiah	cash FX	32	0.505	0.696	0.842	0.368	0.363	0.334	0.331	0.479	0.395
JPY=	US Dollar-Japanese Yen	cash FX	19	0.492	0.797	0.367	0.674	0.568	0.485	0.488	0.665	0.648
KRW=	US Dollar-Korea Won	cash FX	34	0.432	1.204	0.993	0.503	0.431	0.439	0.475	0.372	0.454
MXN=	US Dollar-Mexico Peso	cash FX	24	2.082	1.992	1.206	1.481	1.494	1.416	1.437	1.622	1.507
NOK=	US Dollar-Norwegian Krone	cash FX	33	2.003	1.242	2.095	0.921	0.983	1.103	0.996	0.88	0.981
NZD=	New Zealand Dollar-US Dollar	cash FX	28	1.561	0.597	0.745	0.459	0.488	0.484	0.486	0.512	0.518
PLN=	US Dollar - Polish Zloty	cash FX	39	1.867	0.86	0.538	0.703	0.578	0.544	0.554	0.536	0.703
RUB=	US Dollar - Russian Ruble	cash FX	27	1.178	0.546	1.098	0.382	0.392	0.372	0.371	0.657	0.451
SEK=	US Dollar - Swedish Krone	cash FX	21	0.695	2.987	2.512	0.933	0.95	0.938	1.006	0.908	1.005
SGD=	US Dollar-Singapore Dollar	cash FX	25	0.617	1.808	1.537	0.259	0.272	0.284	0.278	0.268	0.264
TRY=	US Dollar - Turkish Lira	cash FX	24	0.392	2.17	1.312	0.419	0.287	0.32	0.333	0.402	0.405
TWD=	US Dollar - Taiwanese Dollar	cash FX	28	2.369	0.632	2.281	0.488	0.311	0.499	0.501	0.669	0.489
ZAR=	US Dollar - South African Rand	cash FX	40	0.429	0.255	0.355	0.215	0.224	0.239	0.234	0.217	0.222
.BCOM	Bloomberg Commodity Index	commodities	34	0.586	2.245	0.984	0.626	0.425	0.399	0.392	0.481	0.45
ITTEU5Y=MG	ITRAXX Europe CDS index	credit	31	2.132	2	1.871	0.614	0.677	0.608	0.601	0.821	0.668

Table 5: full experiment results, first half

ric	description	asset class	p	ar(1)	kernel ridge	ridge	fte	pwe	rbfnet gmm	rbfnet km	rbfnet rda	ewrls
ITEXO5Y=MG	ITRAXX Crossover CDS index	credit	29	2.549	2.051	2.263	0.611	0.581	0.535	0.527	0.7	0.562
.TRXFLDGLPUENE	Refinitiv Global Energy Index	energy	46	0.682	1.195	0.722	0.42	0.517	0.642	0.659	0.496	0.51
.AEX	AEX Index	equities	50	1.027	0.805	0.957	0.971	0.908	0.898	0.896	0.914	0.928
.AORD	ASX All Ordinaries Index	equities	32	1.406	1.535	1.682	0.92	0.772	0.724	0.719	0.926	0.709
.AXJO	S&P/ASX 200	equities	10	1.063	1.261	1.736	1.288	0.865	0.744	0.725	1.27	0.716
.BFX	BEL 20 Index	equities	49	0.674	0.66	2.055	0.546	0.799	0.706	0.707	1.287	0.448
.FCHI	CAC 40 Index	equities	49	1.17	1.55	0.71	0.811	0.735	0.754	0.749	0.799	0.613
.FTAS	FTSE ALL SHARE	equities	49	1.799	1.15	1.252	1.043	0.822	0.811	0.805	0.887	0.747
.FTSE	FTSE 100 Index	equities	50	0.817	1.74	0.413	1.058	0.89	0.906	0.9	0.913	0.808
.GDAXI	DAX Index	equities	22	1.159	1.157	1.206	0.716	0.737	0.759	0.758	0.704	0.763
.IBEX	IBEX 35 Index	equities	29	0.862	1.346	0.732	0.628	0.603	0.618	0.619	0.728	0.458
.IRTS	RTS Index	equities	50	1.155	1.002	0.265	0.471	0.423	0.38	0.379	0.612	0.299
.IXIC	NASDAQ Composite	equities	11	1.986	0.633	1.206	0.763	0.76	0.749	0.741	0.795	0.756
.MID	S&P 400 Mid Cap Index	equities	26	2.986	2.046	2.481	0.755	0.767	0.773	0.771	0.746	0.781
.NDX	NASDAQ 100	equities	10	1.443	2.174	1.588	0.784	0.741	0.693	0.728	0.79	0.755
.NYA	NYSE Composite	equities	24	1.755	2.875	0.681	0.642	0.649	0.664	0.665	0.669	0.6
.OMXHPI	OMX Helsinki	equities	50	1.779	1.165	1.008	0.456	0.698	0.833	0.84	0.593	0.49
.OMXS30	OMX Stockholm 30 Index	equities	50	2.321	2.343	1.806	1.182	1.018	0.937	0.937	1.24	0.939
.SSMI	Swiss Market Index	equities	22	0.458	1.037	1.17	0.572	0.584	0.612	0.616	0.548	0.575
.STOXX50	EURO STOXX 50 Index	equities	48	1.979	1.6	1.36	0.933	0.793	0.838	0.839	0.818	0.638
.STOXX50E	EURO STOXX 50	equities	49	1.703	0.946	0.913	0.917	0.799	0.808	0.808	0.876	0.674
.TRXFLDGLPU	Refinitiv Global Equities Index	equities	41	1.68	0.838	0.578	0.482	0.609	0.647	0.642	0.693	0.479
.TRXFLDGLPUHLC	Refinitiv Global Healthcare Index	equities	36	0.935	0.914	1.56	0.612	0.73	0.792	0.776	0.806	0.592
AU10YT=RR	Australia 10-year Note	interest rates	43	1.3	1.293	1.189	0.624	0.801	0.616	0.616	1.353	0.659
CH10YT=RR	Swiss 10-year Note	interest rates	28	1.73	0.795	4.464	1.047	0.982	1.022	1.031	0.76	1.238
CN10YT=RR	China 10-year Note	interest rates	20	0.85	0.85	0.985	0.578	0.604	0.623	0.6	0.618	0.577
DE10YT=RR	Germany 10-year Note	interest rates	40	1.74	3.084	2.047	0.903	0.889	0.818	0.848	0.995	0.909
ES10YT=RR	Spain 10-year Note	interest rates	34	1.752	1.789	1.748	0.813	0.82	0.803	0.809	0.864	0.806
FR10YT=RR	France 10-year Note	interest rates	29	1.487	1.592	1.707	0.869	0.852	0.821	0.834	0.91	0.845
GB10YT=RR	United Kingdom 10-year Note	interest rates	40	1.642	2.375	2.216	0.942	0.842	0.726	0.733	1.18	0.811
IT10YT=RR	Italy 10-year Note	interest rates	14	1.295	3.434	3.77	1.351	1.296	1.252	1.258	1.361	1.311
JP10YT=RR	Japan 10-year Note	interest rates	14	1.888	3.129	1.23	0.934	1.22	1.459	1.474	0.819	1.201
RU10YT=RR	Russia 10-year Note	interest rates	12	11.319	6.249	2.829	3.049	2.322	2.036	2.283	2.48	2.568
US10YT=RRPS	10-Year Note	interest rates	7	8.741	12.063	2.408	2.683	2.645	2.589	2.322	2.902	2.795
US2YT=RRPS	2-Year Note	interest rates	17	1.889	4.913	1.471	1.179	1.29	1.44	1.358	1.44	1.096
US30YT=RRPS	30-Year Bond	interest rates	11	0.9	6.086	5.717	0.75	1.089	1.016	1.101	1.108	1.153
US5YT=RRPS	5-Year Note	interest rates	36	1.12	1.284	7.247	0.86	0.838	0.839	0.837	0.887	0.83
ZA10YT=RR	South Africa 10-year Note	interest rates	47	2.139	1.509	0.643	0.571	0.597	0.588	0.584	0.646	0.573
XAG=	Silver	spot metals	23	2.282	2.042	2.275	1.948	1.57	1.256	1.287	1.877	1.973
XAU=	Gold	spot metals	19	2.295	2.169	2.842	1.459	1.419	1.248	1.237	1.571	1.678
XPD=	Palladium	spot metals	25	0.861	2.574	0.818	0.531	0.528	0.531	0.52	0.554	0.532
XPT=	Platinum	spot metals	31	1.855	1.898	2.506	0.64	0.671	0.629	0.681	0.83	0.654

Table 6: full experiment results, second half

Appendix B. Feature Importances

In this section, we show visually, the 10 most selected features of the feature selection algorithm 1. The feature that is selected most, is minutely log returns of .BCOM, the Bloomberg Commodity Index. This is followed by cash silver, the US Dollar vs South African Rand, cash gold, the German DAX index, the Australian Dollar vs Japanese Yen, the Russian RTS index, US Dollar vs Mexican Peso, US Dollar vs Canadian Dollar and US Dollar vs Norwegian Krone. The next plot shows these feature importances by sector. Cash FX instruments provide the greatest contribution, selected over 40% of the time. Equities provides the second highest contribution, at over 25%.

Figure 4: ranked feature importances

Figure 5: feature importances by sector

Appendix C. Model Structure versus Performance

The plot and tables that follow, shows model structure versus normalised mean-square error, for the radial basis function networks formulated out the of kmeans++ and Gaussian mixture models of section 3.2, as well as the exponentially weighted recursive least-squares model of section 3.4. For the exponentially weighted recursive least-squares model, the variable k indicates the number of features p selected by the meta-algorithm 2.3. For the radial basis function networks, k indicates the number of hidden processing units chosen by their respective clustering algorithms. We see that meta-algorithm 2.3 selects on average $28/82 \approx 34\%$ of available predictors. In addition, the average number of hidden processing units for the radial basis function network formulated out of Gaussian mixture models averages 6, whereas the ones formulated out of kmeans++, averages 3. The effective number of parameters to be estimated for the *rbfnet gmm* is thus on average 28×6 means + $28^2 \times 6$ covariances matrices + a 29 element weight vector for the supervised learner, which totals 4901 parameters. Table 8 shows that the normalised mean-square prediction error decreases for the exponentially weighted recursive least-squares model, as the number of predictors selected by the feature selection meta-algorithm increases. This is measured through correlation. It also shows that the normalised mean-square prediction error tends to decrease, albeit at a slower rate, as the number of radial basis function hidden processing units decrease.

Figure 6: model structure versus performance

		ewrls	rbfnet gmm	rbfnet km
k	count	1230	1230	1230
k	mean	28	6	3
k	std	12	5	1
k	min	7	3	2
k	max	50	20	11
nmspe test	count	1230	1230	1230
nmspe test	mean	0.87	0.834	0.833
nmspe test	std	0.56	0.459	0.461
nmspe test	min	0.221	0.237	0.232
nmspe test	max	3.039	2.59	2.323

Table 7: model structure versus performance, summary

model	k vs nmspe test
ewrls	-0.365
rbfnet gmm	0.121
rbfnet km	0.05

Table 8: model structure versus performance, correlation matrix

References

- B. Abraham and J. Ledolter. *Statistical methods for forecasting*. Wiley, 1983.
- D. Arthur and S. Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, page 1027–1035, USA, 2007. Society for Industrial and Applied Mathematics. ISBN 9780898716245.
- D. Barber, C. Taylan, and C. Silvia. *Bayesian Time Series Models*. Cambridge University Press, 2011.
- S. A. Billings, C. F. Fung, , and W. Luo. *On-Line Supervised Adaptive Training Using Radial Basis Function Networks*. Department of Automatic Control and Systems Engineering, 1996.
- C. M. Bishop. *Neural networks for pattern recognition*. Oxford University Press, 1995.
- F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- G. Bontempi. Long term time series prediction with multi-input multi-output local learning. *Proceedings of the 2nd European Symposium on Time Series Prediction (TSP), ESTSP08*, 01 2008.
- D. Broomhead and D. Lowe. Multivariable functional interpolation and adaptive networks. *Complex Syst.*, 2, 1988.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006. ISBN 9780511546921.
- C. Chatfield. *The analysis of time series : an introduction with R*. Chapman & Hall, CRC, seventh edition edition, 2019.
- S. Chen and S. A. Billings. Representations of non-linear systems: the narmax model. *International Journal of Control*, 49(3):1013–1032, 1989. doi: 10.1080/00207178908559683.
- N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press, 2000. doi: 10.1017/CBO9780511801389.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39:1–38, 1977. ISSN 00359246.
- D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366):427, 1979.
- R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. Wiley, second edition. edition, 2001. ISBN 9781118586006.
- C. Engel. Can the markov switching model forecast exchange rates? *Journal of international economics*, 36(1):151–165, 1994. ISSN 0022-1996.
- M. A. T Figueiredo and A. K Jain. Unsupervised learning of finite mixture models. *IEEE transactions on pattern analysis and machine intelligence*, 24:381–396, 2002. ISSN 0162-8828. doi: 10.1109/34.990138.
- J. H. Friedman. Regularized discriminant analysis. *Journal of the American Statistical Association*, 84(405):165–175, 1989. ISSN 01621459.

- F. Girosi and T. Poggio. Networks and the best approximation property. *Biological cybernetics*, 63(3):169–176, 1990.
- C.W.J. Granger and P. Newbold. Spurious regressions in econometrics. *Journal of econometrics*, 2:111–120, 1974.
- S. Gunnarsson. Combining tracking and regularization in recursive least squares identification. In *Proceedings of 35th IEEE Conference on Decision and Control*, volume 3, pages 2551–2552 vol.3, 1996. doi: 10.1109/CDC.1996.573481.
- R. Hadsell, D. Rao, A. Rusu, and R. Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24:1028–1040, 2020.
- A. Harvey. *Time series models*. Financial Times/Prentice Hall, 2nd ed. edition, 1993.
- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning : data mining, inference, and prediction*. Springer series in statistics. Springer Verlag, second edition. edition, 2009.
- T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC, 2015.
- A. Hoerl. Application of ridge analysis to regression problems. In *Chemical Engineering Progress*, 1962.
- H. Ince, A. F. Cebeci, and S. Z. Imamoglu. An artificial neural network-based approach to the monetary model of exchange rate. *Computational economics*, 53(2):817–831, 2019. ISSN 0927-7099.
- A. S. Iwashita and J. P. Papa. An overview on concept drift learning. *IEEE access*, 7:1532–1547, 2019.
- R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35, 1960.
- N. Kanazawa. Radial basis functions neural networks for nonlinear time series analysis and time-varying effects of supply shocks. *Journal of macroeconomics*, 64:103210, 2020.
- C. Kaplanis, M. Shanahan, and C. Clopath. Continual reinforcement learning with complex synapses, 2018.
- H. Khosravi. A novel structure for radial basis function networks—wrbf. *Neural processing letters*, 35(2):177–186, 2011.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes, 2014.
- J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114:3521–3526, 2017.
- A. Koshiyama, S. B. Blumberg, S. Flennerhag, N. Firoozye, and P. Treleaven. Quantnet: Transferring learning across systematic trading strategies, 2020.
- W. Liu, J. C. Principe, and S. Haykin. *Kernel adaptive filtering a comprehensive introduction*. Wiley, 2010.

- S. Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2): 129–137, 1982.
- D. J.C. MacKay. Bayesian non-linear modelling for the prediction competition. In *In ASHRAE Transactions*, V.100, Pt.2, pages 1053–1062. ASHRAE, 1994.
- I. Markovsky and S. Van Huffel. Overview of total least-squares methods. *Signal Processing*, 87(10): 2283–2302, 2007. ISSN 0165-1684. Special Section: Total Least Squares and Errors-in-Variables Modeling.
- R. A. Meese and K. Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics*, 14(1):3–24, 1983. ISSN 0022-1996.
- R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1):125 – 144, 1976.
- J. Moody and C. J. Darken. Fast learning in networks of locally-tuned processing units. *Neural computation*, 1:281–294, 1989.
- E. Moroshko, N. Vaits, and K. Crammer. Second-order non-stationary online learning for regression. *Journal of Machine Learning Research*, 16(43):1481–1517, 2015.
- K. P. Murphy. *Machine learning a probabilistic perspective*. Adaptive Computation and Machine Learning. MIT Press, 2012.
- R. M. Neal. *Bayesian learning for neural networks / by Radford M. Neal*. Lecture notes in statistics (Springer-Verlag) ; 118. Springer, New York, 1996. ISBN 9781461207450.
- C. E. Rasmussen and J. Quiñonero-candela. Healing the relevance vector machine through augmentation. In *In Proc. of ICML 22*, 2005.
- P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987. ISSN 0377-0427.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0.
- G. A. Rummery and M. Niranjan. On-line q-learning using connectionist systems. In *Technical Report CUED/F-INFENG/TR 166*, Cambridge University Engineering Department, 1994.
- B. K. Salvalaio and G. de Oliveira Ramos. Self-adaptive appearance-based eye-tracking with online transfer learning. In *2019 8th Brazilian Conference on Intelligent Systems (BRACIS)*, pages 383–388. IEEE, 2019.
- J. Sherman and W. J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Annals of Mathematical Statistics*, 21:124–127, 1950.
- S. Sukhov, M. Leontev, A. Miheev, and K. Sviatov. Prevention of catastrophic interference and imposing active forgetting with generative methods. *Neurocomputing (Amsterdam)*, 400:73–85, 2020.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B, Methodological*, 58(1):267–288, 1996.

- M. E. Tipping. The relevance vector machine. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, NIPS'99, page 652–658, Cambridge, MA, USA, 1999. MIT Press.
- V. N. Vapnik. *Statistical learning theory*. Adaptive and learning systems for signal processing, communications, and control. Wiley, New York ; Chichester, 1998. ISBN 0471030031.
- P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(110):3371–3408, 2010.
- V. Vovk. *Kernel Ridge Regression*, pages 105–116. Springer, 2013.
- C. S. Wallace and D. L. Dowe. Minimum Message Length and Kolmogorov Complexity. *The Computer Journal*, 42(4):270–283, 01 1999. ISSN 0010-4620. doi: 10.1093/comjnl/42.4.270.
- J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumanan, and M. Botvinick. Learning to reinforcement learn, 2017.
- X. Wang, X. Wang, and Z. Zeng. A novel weight update rule of online transfer learning. In *2020 12th International Conference on Advanced Computational Intelligence (ICACI)*, pages 349–355. IEEE, 2020.
- C. J. Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, Oxford, 1989.
- C. Williams and M. Seeger. Using the nystrom method to speed up kernel machines. In T. Leen, T. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems*, volume 13. MIT Press, 2001.
- Q. Yang, Y. Zhang, W. Dai, and S. J. Pan. *Transfer Learning*. Cambridge University Press, 2020.
- P. Zhao, S. C. H. Hoi, J. Wang, and B. Li. Online transfer learning. *Artificial intelligence*, 216: 76–102, 2014.