
Recurrent Neural Networks and Universal Approximation of Bayesian Filters

Adrian N. Bishop

University of Technology Sydney, and CSIRO

Edwin V. Bonilla

CSIRO's Data61

Abstract

We consider the Bayesian optimal filtering problem: i.e. estimating some conditional statistics of a latent time-series signal from an observation sequence. Classical approaches often rely on the use of assumed or estimated transition and observation models. Instead, we formulate a generic recurrent neural network framework and seek to learn directly a recursive mapping from observational inputs to the desired estimator statistics. The main focus of this article is the approximation capabilities of this framework. We provide approximation error bounds for filtering in general non-compact domains. We also consider strong time-uniform approximation error bounds that guarantee good long-time performance. We discuss and illustrate a number of practical concerns and implications of these results.

1 INTRODUCTION

Optimal filtering (Bain and Crisan, 2008) is concerned with estimating some statistics of a latent random signal (or state) X_t at the current time $t \in \mathbb{N}$, conditioned on some observations Y_τ collected thus far, i.e., $0 \leq \tau \leq t$. When the signal transition and observation models are linear with additive Gaussian noise, the solution is given by the celebrated Kalman filter (Kalman, 1960). In general non-linear, non-Gaussian, settings, there is no tractable finite-dimensional optimal filter and approximations are needed. Nonlinearity and high state-dimensionality, typical in applications, make the filtering problem challenging.

We consider a data-driven approach to learning an optimal estimator. More specifically, we consider scenarios in which latent signal and observation data may be collected by computer simulations (if a good model is available), actual experiments (otherwise) or some combination thereof.

We provide a framework for optimal filtering using general recurrent neural net (RNN) structures that may be trained by minimizing the empirical mean-square error between the sampled signal data and the network output, with the network taking as input only the current observation at each time. The output of the RNN acts as a recursive state estimator. We study a number of interesting approximation capabilities of this estimator with respect to the true optimal nonlinear filter (e.g. the true conditional expectations).

Summary of contributions: We propose a generic RNN-based architecture and methodology for optimal (Bayesian) filtering in general state-space models. This formulation is kept generic so as to facilitate a study of the approximation capabilities and limitations of RNNs in applications to Bayesian filtering. Firstly: We show that a generic RNN-based estimator can approximate the optimal estimate of the signal to any desired accuracy on a fixed finite time interval of interest. We note that almost no assumptions on the state-space model are needed in this case. Furthermore, the model may not even be known if experimental data is available to generate signal training data. However, the approximation capability in this setting relies strictly on the assumption of a finite horizon of interest. Secondly: We study the time-uniform universal approximation of Bayesian filters with RNNs. For a particular class of models, we show the approximation error can be bounded, to any desired accuracy, uniformly *for all time*; i.e. implying that approximation errors do not accumulate over time. This result has important practical implications, e.g. it may influence network design, and it allows one to train on a signal-observation sequence of (short) finite length while permitting the filtering algorithm to run indefinitely (as is typical in applications). Thirdly: We discuss and illustrate a number of practical consequences of both results and contrast these with each other and with other methodologies.

Background and related work in filtering: Optimal filtering and related problems in learning and inference in dynamical systems are of interest across many fields of study, including control and signal processing (Anderson and Moore, 1979; Bain and Crisan, 2008), geophysics Evensen (2009), machine learning (Van Der Merwe et al., 2000; Ghahramani and Jordan, 1995), and statistics (Pitt and Shephard, 1999; Andrieu et al., 2010). A comprehen-

sive review of the methodology in this area is beyond the scope of this paper. We simply note in passing some popular model-based approximations for nonlinear filtering such as the extended Kalman filter (EKF) (Anderson and Moore, 1979) and the unscented Kalman filter (UKF) (Julier and Uhlmann, 2004), and Monte Carlo integration methods for filtering, termed particle filters (Gordon et al., 1993; Pitt and Shephard, 1999), and ensemble Kalman filter (ENKF) (Evensen, 2009) methods. There are some adaptations of these model-based methods to data-driven, likelihood-free, filtering, e.g., employing approximate Bayesian computation within Monte Carlo (Jasra et al., 2010; Martin et al., 2019) or using Gaussian processes (Ko and Fox, 2009).

Neural network approaches to optimal nonlinear filtering were considered in early work in Lo (1994); Parisini and Zoppoli (1994); V. T. and Shin (1994); Alessandri et al. (1999); Parlos et al. (2001). The article Haykin et al. (1997) provides a good early summary and survey. These works naturally considered very simple networks, e.g. one layer and just a handful of sigmoidal-type activation functions; e.g. see Lo (1994). However, convergence and approximation error results were also given in Lo (1994); Parisini and Zoppoli (1994) and also in Parisini et al. (1997); Alessandri et al. (1999). We generalize, strengthen, and add to these early results in a number of ways in this work; e.g. in particular with our main result on time-uniform approximation bounds with observations on non-compact domains.

Other work has used machine learning to improve the optimization of calculating the initial state in a variational data assimilation framework (Frerix et al., 2021) and inference in state-space models (see, e.g., Murphy, 2012, Ch. 18). Some approaches have developed approximate inference techniques in parametric settings (Ghahramani and Hinton, 2000; Fox et al., 2008). Later methods have used non-parametric models such as Gaussian processes (Frigola et al., 2014; Doerr et al., 2018; Nickisch et al., 2018; Ialongo et al., 2019) or flexible modern neural network-based frameworks (Krishnan et al., 2017; Gu et al., 2015; Karl et al., 2016; Haarnoja et al., 2016; Becker et al., 2019). Finally, we note that learning time-series models with neural networks is closely related to the filtering discussed here; e.g., it may be considered a special case in which the observed data is just the signal process with no noise, see, e.g., Karl et al. (2016); Rangapuram et al. (2018).

The purpose of this work is not the development of new (RNN-based) methodology for Bayesian filtering. Instead, we study the universal approximation capability of (rather generic) RNN-based approximations of the optimal filter, both on finite time intervals and uniformly in time. In later sections, when we can easily reference specific technical details and conditions, we discuss the results of this article, and (briefly) contrast these with other approximation methods as in Heine and Crisan (2008); van Handel (2009); Whiteley (2013); Douc et al. (2014); Crisan et al. (2020).

2 DISCRETE-TIME BAYESIAN FILTERING

Let $\mathbb{X} \subseteq \mathbb{R}^{d_x}$ and $\mathbb{Y} \subseteq \mathbb{R}^{d_y}$ with the Borel σ -algebra $\mathcal{B}(\mathbb{X})$. Consider a Markov chain $(X_t)_{t \in \mathbb{N}}$ taking values in \mathbb{X} with Markov kernel $\mathsf{K} : \mathbb{X} \times \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$. Consider a process $(Y_t)_{t \in \mathbb{N}}$ defined on \mathbb{Y} , conditionally independent given $(X_t)_{t \in \mathbb{N}}$, with a transition density $g : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ with respect to the Lebesgue measure. The process X_t is thought of as being observed via the process Y_t . The process $(Y_t)_{t \in \mathbb{N}}$ is itself not a Markov chain, but the pair $(X_t, Y_t)_{t \in \mathbb{N}}$ is a Markov chain on $\mathbb{X} \times \mathbb{Y}$. Denote¹ by \mathbf{P}^μ the law of $(X_t, Y_t)_{t \in \mathbb{N}}$ under which the pair $(X_t, Y_t)_{t \in \mathbb{N}}$ is a Markov chain on $\mathbb{X} \times \mathbb{Y}$ with X_0 having measure $\mu \in \mathcal{P}(\mathbb{X})$. The space of all probability measures on \mathbb{X} is $\mathcal{P}(\mathbb{X})$. Expectation with respect to \mathbf{P}^μ is denoted by \mathbf{E}^μ .

Filtering involves computing the regular conditional distribution $\pi_t^\mu(A) := \mathbf{P}^\mu(X_t \in A | Y_1, \dots, Y_t)$, for all $A \in \mathcal{B}(\mathbb{X})$. The distribution $\pi_t^\mu \in \mathcal{P}(\mathbb{X})$ is called the filtering distribution. From Bayes' rule,

$$\begin{aligned} \pi_0^\mu &= \mu, \\ \pi_t^\mu &= \Psi(\pi_{t-1}^\mu, Y_t) \end{aligned} \quad (2.1)$$

where, for some $\nu \in \mathcal{P}(\mathbb{X})$ and $y \in \mathbb{Y}$, the function $\Psi : \mathcal{P}(\mathbb{X}) \times \mathbb{Y} \rightarrow \mathcal{P}(\mathbb{X})$ is defined by,

$$\Psi(\nu, y)(A) := \frac{\int_A g(x, y) (\nu \mathsf{K})(dx)}{\int g(x, y) (\nu \mathsf{K})(dx)}, \quad \forall A \in \mathcal{B}(\mathbb{X}) \quad (2.2)$$

where $(\nu \mathsf{K})(dx) := \int \mathsf{K}(x', dx) \nu(dx')$. The mapping Ψ is time-invariant. The process $(\pi_t^\mu)_{t \in \mathbb{N}}$ on $\mathcal{P}(\mathbb{X})$ is a Markov process when X_0 has measure $\mu \in \mathcal{P}(\mathbb{X})$ (Stettner, 1989).

For a finite-dimensional integrable function $x \mapsto \rho(x)$, $x \in \mathbb{X}$, taking values in some Euclidean space, the filtering problem is often stated in terms of a point-valued estimate,

$$\pi_t^\mu(\rho) := \int \rho(x) \pi_t^\mu(dx) = \mathbf{E}^\mu[\rho(X_t) | Y_1, \dots, Y_t] \quad (2.3)$$

For example, one may want to estimate just the conditional mean and covariance of X_t , if they exist.

2.1 Bayesian Filtering as Optimal State Estimation

A closely related problem to that of filtering is state estimation. Going forward, we assume (X_t, Y_t) are jointly square integrable and any measurable function preserves square integrability, e.g. $\mathbb{E}[\|\rho(X_t)\|^2] < \infty$. Then, consider the measurable function $\bar{\rho}_t^\mu := \bar{\rho}_t^\mu(Y_1, \dots, Y_t)$; being the solution of the optimization problem,

$$\bar{\rho}_t^\mu := \operatorname{argmin}_f \mathbf{E}^\mu [\|f(Y_1, \dots, Y_t) - \rho(X_t)\|^2] \quad (2.4)$$

¹Superscripts often denote initial conditions for random processes, but are also indices (e.g. over data, parameters, etc), and also just powers in some places. The use and case should be clear.

where f is any measurable function of (Y_1, \dots, Y_t) . The solution to (2.4) is related to the optimal filtering problem by $\bar{\rho}_t^\mu = \pi_t^\mu(\rho)$, as defined in (2.3). Note $\bar{\rho}_t^\mu$ is square integrable (as it may be viewed as the orthogonal projection onto the closed subspace of square integrable functions).

We think of $\bar{\rho}_t^\mu$ as the optimal estimate of the state $\rho(X_t)$. This estimate is intractable in all but the most pathological situations. It is desirable in practice to design an approximation of $\bar{\rho}_t^\mu$. We denote a computable estimate by $\hat{\rho}_t^\nu$ where $\nu \in \mathcal{P}(\mathbb{X})$ is a known, assumed, or estimated distribution of X_0 such that $\hat{\rho}_0^\nu := \nu(\rho)$ is also computable.

We emphasise here that the solution to (2.4), as given by $\bar{\rho}_t^\mu = \pi_t^\mu(\rho)$, see (2.3), is not generally given in the form of a recursion, even though the recursion (2.1) on the level of the conditional distribution exists.

To close this section, note $\pi_t^\mu(\rho)$ makes sense in more general settings than outlined for (X_t, Y_t) thus far (e.g. beyond independent, time-invariant, Markov settings); as does the solution to (2.4). Recursions of the form (2.1) also hold under more general models, see [Tong and van Handel \(2012\)](#).

2.2 Sufficient Coordinates and Recursive Filtering

The recursion in (2.1) is on the level of probability measures and is generally infinite dimensional and intractable. Informally, e.g. neglecting existence and other structure, we may define a (generally infinite dimensional) vector state of sufficient statistics on a (separable) Banach space \mathbb{S} . For the filtering distribution $\pi_t^\mu \in \mathcal{P}(\mathbb{X})$ we may consider, for example, the conditional moments,

$$S_t^\mu := \int \begin{bmatrix} \text{vec}(xx^\top) \\ \vdots \end{bmatrix} \pi_t^\mu(dx) \quad (2.5)$$

as in [Rudenko \(2010\)](#). A recursive Bayes filter may then take the form,

$$S_t^\mu = \Psi(S_{t-1}^\mu, Y_t) =: \Psi_t(S_{t-1}^\mu) \quad (2.6)$$

on \mathbb{S} for some measurable function $\Psi : \mathbb{S} \times \mathbb{Y} \rightarrow \mathbb{S}$. The subscript t in Ψ_t succinctly indexes the observational input; highlighting the recursion on \mathbb{S} . In this notation, we write the composition $S_t^\mu = \Psi_t \Psi_{t-1} \dots \Psi_{\tau+1}(S_\tau^\mu)$ for $\tau < t \in \mathbb{N}$. If such a representation exists, then Ψ is time-invariant.

We may claim that a finite-dimensional recursive optimal filter exists if and only if \mathbb{S} is finite dimensional and we can transform between S_t^μ and π_t^μ in some standard way. In the linear-Gaussian setting a finite-dimensional state S_t^μ and recursion (2.6) exists from which π_t^μ follows ([Anderson and Moore, 1979](#)). See also [Sawitzki \(1981\)](#); [Dickinson and Sontag \(1985\)](#); [Levine and Pignie \(1986\)](#). If S_t^μ is finite, then it is a special case of $\pi_t^\mu(\rho)$ for some $x \mapsto \rho(x)$.

3 GENERIC RNN-BASED STATE ESTIMATORS

We introduce a generic RNN architecture for state estimation. Consider a network with $L \in \mathbb{N}$ layers, in the form,

$$\hat{\rho}_t^\nu = \mathbf{W}_{(L,L-1)} \mathbf{s}_t^{(L-1)} + \mathbf{b}_{(L)} \quad (3.1)$$

$$\mathbf{s}_t^{(l)} = \sigma \left(\mathbf{W}_{(l,l-1)} \mathbf{s}_t^{(l-1)} + \mathbf{b}_{(l)} + \sum_{k=l}^{L-1} \mathbf{W}_{(l,k)} \mathbf{s}_{t-1}^{(k)} \right), \quad l \in \{1, \dots, L-1\} \quad (3.2)$$

$$\mathbf{s}_t^{(0)} := Y_t \quad (3.3)$$

where $\sigma(\cdot)$ is a vector-valued activation function acting component-wise on its argument. Component-wise, i.e. on scalar inputs, we assume $\sigma(\cdot) := \max\{0, \cdot\}$ here, i.e. we consider only rectified linear activation units (RELUs). The parameter $\theta := [(\mathbf{W}_{(l,l-1)}, \mathbf{W}_{(l,k)}), \mathbf{b}_{(l)}]$ consists of appropriately sized real matrices and vectors. The generic notation for the RNN-approximated state estimator is,

$$\hat{\rho}_t^\nu = \hat{\Psi}^\theta(\mathbf{s}_{t-1}, Y_t) := \hat{\Psi}_t^\theta(\mathbf{s}_{t-1}) \quad (3.4)$$

where \mathbf{s}_t is a stacked vector of the $\mathbf{s}_t^{(l)}$. The parameter θ is time-invariant and takes values in a high-dimensional Euclidean space. The recursion (3.4) is on the level of the internal network state \mathbf{s}_{t-1} which we suggest captures in some sense a finite-dimensional approximation of sufficient statistics as in (2.6), enough to approximate $\pi_t^\mu(\rho)$.

Assumption 1. *The process $(X_t, Y_t)_{t \in \mathbb{N}}$ is jointly square integrable and $x \mapsto \rho(x)$, and $y \mapsto \hat{\Psi}^\theta(\cdot, y)$, are finite-dimensional, square integrable, and take values in some Euclidean space. The assumption on $\hat{\Psi}^\theta$ is provable.*

Common specialisations of (3.1), (3.2), (3.3) involve feedback solely from the last activation layer to the first layer; and feedback from the output of each layer to itself, see [Pascanu et al. \(2014\)](#). Our results are proven under common simplifications of (3.1), (3.2), (3.3).

The ideal RNN-approximated estimator is based on optimizing θ according to the following cost functional,

$$\mathcal{C}(\theta) := \frac{1}{T} \sum_{t=1}^T \mathbf{E}^\mu \left[\|\hat{\Psi}^\theta(\mathbf{s}_{t-1}, Y_t) - \rho(X_t)\|^2 \right] \quad (3.5)$$

on a finite horizon $T \in \mathbb{N}$. In practice, suppose some data $\mathcal{D}_{T,N} := (X_t^{(n)}, Y_t^{(n)})$, $t \in \{1, \dots, T\}$, $n \in \{1, \dots, N\}$ with finite $N \in \mathbb{N}$, is independently sampled according to the law \mathbf{P}^μ . The data may be collected via computer simulations (if a good model \mathbf{P}^μ is available), real-world experiments, or some combination. Practically, a network is trained by minimizing, over θ , the empirical loss,

$$\theta^* := \underset{\theta}{\operatorname{argmin}} \mathcal{C}^N(\theta), \quad (3.6)$$

$$\mathcal{C}^N(\theta) := \frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T \|\hat{\Psi}^\theta(\mathbf{s}_{t-1}, Y_t^{(n)}) - \rho(X_t^{(n)})\|^2$$

with \mathbf{s}_0 given. The loss $\mathcal{C}^N(\theta)$ is a sample version of (3.5) and $\mathcal{C}^N(\theta) \xrightarrow{N \rightarrow \infty} \mathcal{C}(\theta)$. In practice, we might define $\hat{\rho}_0^\nu := \nu(\rho)$ and set \mathbf{s}_0 commensurately, dependent on the network. We may also consider \mathbf{s}_0 as a model parameter and optimise its value (Lo, 1994). The approximation theorems consider the initialisation as part of the result.

4 APPROXIMATION THEOREMS

4.1 A General Approximation Theorem

Theorem 1. *Let Assumption 1 hold. Let $\hat{\Psi}^\theta(\mathbf{s}_{t-1}, Y_t)$ denote a generic multilayer RNN, as in (3.1), (3.2), (3.3), taking as input elements in the sequence $(Y_t)_{t \in \{1, \dots, T\}}$. Then for any $\epsilon > 0$, there exists a finite real parameter vector θ^* , and an initialisation vector \mathbf{s}_0 , such that,*

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E}^\mu \left[\|\pi_t^\mu(\rho) - \hat{\Psi}^{\theta^*}(\mathbf{s}_{t-1}, Y_t)\|^2 \right]^{\frac{1}{2}} \leq \epsilon \quad (4.1)$$

Proof. Note $\pi_t^\mu(\rho) = \mathbb{E}[\rho(X_t) | Y_1, \dots, Y_t]$. This is an extension of Lo (1994), who considers single-layer networks with observations on compact domains. Instead we use Kidger and Lyons (2020) and consider deep RNNs on non-compact domains. We construct a network of the form,

$$\hat{\rho}_t^\nu = \mathbf{W}_{(L, L-1)} \mathbf{s}_t^{(L-1)} + \mathbf{b}_{(L)} \quad (4.2)$$

$$\mathbf{s}_t^{(l)} = \sigma(\mathbf{W}_{(l, l-1)} \mathbf{s}_t^{(l-1)} + \mathbf{b}_{(l)}), \quad l \in \{2, \dots, L-1\} \quad (4.3)$$

$$\mathbf{s}_t^{(1)} = \sigma(\mathbf{W}_{(1,0)} Y_t + \mathbf{W}_{(1,1)} \mathbf{s}_{t-1}^{(1)} + \mathbf{b}_{(1)}) \quad (4.4)$$

with $\mathbf{W}_{(2,1)} = \mathbf{W}_{(2,1.5)} \mathbf{W}_{(1.5,1)}$ and $\mathbf{b}_{(2)} = \mathbf{b}'_{(2)} + \mathbf{b}_{(1.5)}$. Consider the first layer and a half,

$$\mathbf{s}_t^{(1.5)} = \mathbf{W}_{(1.5,1)} \mathbf{s}_t^{(1)} + \mathbf{b}_{(1.5)} \quad (4.5)$$

with $\mathbf{s}_t^{(1.5)}$ being the effective input then to layer 2. We seek at any time $t \in \{1, \dots, T\}$ the state,

$$\mathbf{s}_t^{(1.5)} = \left[t \bar{\rho}_0^\mu \top Y_t \top Y_{t-1} \top \dots Y_1 \top 0 \ 0 \ \dots \ 0 \right]^\top \quad (4.6)$$

with $\mathbf{s}_T^{(1.5)} = [T \bar{\rho}_0^\mu \top Y_T \top \dots Y_1 \top]$. Let $\mathbf{W}_{(1.5,1)} = \mathbf{I}$ and,

$$\mathbf{W}_{(1,0)} = \begin{bmatrix} [0 \ 0 \ \dots \ 0] \\ \mathbf{0} \\ \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}, \quad (4.7)$$

$$\mathbf{W}_{(1,1)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & 0 & 0 & 0 & \ddots & \ddots & \ddots & 0 & 0 & 0 & 0 \end{bmatrix} \quad (4.8)$$

where \mathbf{I} in $\mathbf{W}_{(1,0)}$ and all but the first \mathbf{I} in $\mathbf{W}_{(1,1)}$ denotes a d_y -dimensional identity matrix, and the first \mathbf{I} in $\mathbf{W}_{(1,1)}$

is the size of $\bar{\rho}_0^\mu$. And let,

$$\mathbf{b}_1 = [1 \ [0 \ \dots \ 0]^\top \ [b \ \dots \ b]^\top \ 0 \ \dots \ 0]^\top, \quad (4.9)$$

$$\mathbf{b}_{(1.5)} = [0 \ -b \ -b \ \dots \ -b]^\top \quad (4.10)$$

Initialise,

$$\mathbf{s}_0 = \mathbf{s}_0^{(1)} = [0 \ [\bar{\rho}_0^\mu + [b \ \dots \ b]^\top]^\top \ [0 \ \dots \ 0] \ b \ \dots \ b]^\top \quad (4.11)$$

With this network construct and with $b > 0$ large enough we find $\forall t \in \{1, \dots, T\}$ the state $\mathbf{s}_t^{(1.5)}$ is given by (4.6) over a domain on which \mathbf{P}^μ places most mass, as desired, see Kidger and Lyons (2020, proof of Theorem 4.16).

There is no more feedback in the network constructed in this proof and the state $\mathbf{s}_t^{(1.5)}$ can thus be viewed as an input for the feedforward neural network from layers 2 to L .

We define a target function, for the feedforward network from layers 2 to L , by,

$$f(t, \bar{\rho}_0^\mu, Y_1, \dots, Y_T) = \begin{cases} \bar{\rho}_0^\mu & \text{if } t = 0 \\ \mathbf{E}^\mu[\rho(X_1) | Y_1] & \text{if } t = 1 \\ \vdots & \\ \mathbf{E}^\mu[\rho(X_T) | Y_T, \dots, Y_1] & \text{if } t = T \end{cases} \quad (4.12)$$

This function is (Borel) measurable. We now apply classical universal approximation theorems on f . See Kidger and Lyons (2020, proof of Theorem 4.16) for an easy to follow construction immediately applicable here. \square

The preceding theorem holds only on those finite time horizons $t \in \{1, \dots, T\}$, owing to the special structure of the constructed network in the proof. The finite time $T \in \mathbb{N}$ may be arbitrary, but the size of the resulting network (in theory) grows with T . We discuss further the results of this work later. Next we show that, for a particular class of models, the approximating network size is not generally a function of time and the approximation error remains bounded uniformly for all time; i.e. implying that approximation errors do not accumulate over time.

4.2 Recursive Filters, Approximations, and Time-Uniform Approximation Error Bounds

The main result is presented in this subsection. We consider a finite-dimensional statistic of interest $\pi_t^\mu(\rho) =: S_t^\mu$ that evolves recursively as per (2.6). We may consider a finite-dimensional truncation of an infinite-dimensional sufficient statistic in place of S_t^μ if necessary. See Sawitzki (1981); Dickinson and Sontag (1985) for general finite-dimensional filtering results and Goodman (1975); Segall (1976); Marcus (1979); Daum (1986) for models, which may lead naturally to finite-dimensional truncated representations of statistics. We need some assumptions.

Assumption 2. *Observation process $(Y_t)_{t \in \mathbb{N}}$ is stationary.*

Assumption 3. *The map $s \mapsto \Psi(s, \cdot)$ is Lipschitz with finite Lipschitz constant, and there exist real finite constants $C > 0$, $0 < \kappa < 1$ independent of $\tau, t \in \mathbb{N}$ such that,*

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| \Psi_t \cdots \Psi_\tau(S_{\tau-1}^{s_0}) - \Psi_t \cdots \Psi_\tau(S_{\tau-1}^s) \right\|^2 \right]^{\frac{1}{2}} \\ \leq C \kappa^{(t-\tau)} \mathbf{E}^\mu \left[\left\| S_{\tau-1}^{s_0} - S_{\tau-1}^s \right\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.13)$$

for any initial points $s_0, s \in \mathbb{S}$ and all $\tau \leq t$.

This (4.13) is an ‘‘eventually (exponentially) contracting on average’’-type property of the optimal filter (2.6). With (4.13), trajectories eventually come together on average, at an exponential rate; but they may never converge to a fixed point or invariant measure. The latent signal need not be ergodic nor admit an invariant measure for this condition to hold; but under general conditions if the signal is (exponentially) ergodic, then one expects so is the Bayesian filter (2.1) or (2.6). Results of this type (but often not as strong as the assumption) have been widely studied, and we reference: Atar and Zeitouni (1997); Budhiraja and Ocone (1999); Chigansky and Liptser (2004); Oudjane and Rubenthaler (2005); Kleptsyna and Veretennikov (2007); Crisan and Heine (2008); van Handel (2008); Tong and van Handel (2012); Gerber and Whiteley (2017). Although the details differ and are challenging, known results support the intuition that sufficiently informative observations or sufficiently contractive latent signals may generally translate to results like in Assumption 3. This idea is well exemplified with linear-Gaussian models, discussed later, see Bishop and Del Moral (2017). The following is the main result.

Theorem 2. *Let $\widehat{\Psi}^\theta(s_{t-1}, Y_t)$ denote a generic multilayer RNN, as in (3.1), (3.2), (3.3), taking as input elements in the sequence $(Y_t)_{t \in \mathbb{N}}$. Suppose Assumptions 1, 2 and 3 hold. Then for any $\epsilon > 0$, there is a finite real parameter θ^* , and an initialisation vector s_0 , such that,*

$$\mathbf{E}^\mu \left[\left\| S_t^{s_0} - \widehat{\Psi}^{\theta^*}(s_{t-1}, Y_t) \right\|^2 \right]^{\frac{1}{2}} \leq \epsilon \quad (4.14)$$

for all $s_0 \in \mathbb{S}$ and all $t \in \mathbb{N}$. This is a time-uniform bound.

Proof. We construct a RNN in the special form,

$$\widehat{S}_t^{s_0} = \mathbf{W}_{(L,L-1)} \mathbf{s}_t^{(L-1)} + \mathbf{b}_{(L)} \quad (4.15)$$

$$\begin{aligned} \mathbf{s}_t^{(l)} &= \sigma(\mathbf{W}_{(l,l-1)} \mathbf{s}_t^{(l-1)} + \mathbf{b}_{(l)}), \\ & \quad l \in \{2, \dots, L-1\} \end{aligned} \quad (4.16)$$

$$\mathbf{s}_t^{(1)} = \sigma(\mathbf{W}_{(1,0)} Y_t + \mathbf{W}_{(1,L-1)} \mathbf{s}_{t-1}^{(L-1)} + \mathbf{b}_{(1)}) \quad (4.17)$$

Only the output of the last activation layer is fed back to the input of the first layer. With no loss of generality, we write,

$$\widehat{S}_t^{s_0} = \widehat{\Psi}_t(\mathbf{s}_{t-1}) = \widehat{\Psi}_t(\widehat{S}_{t-1}^{s_0}) \quad (4.18)$$

The notation $\widehat{S}_t = \widehat{\Psi}_t(\widehat{S}_{t-1})$ and the network structure highlights the recursive nature of this approximation, in line with (2.6). This feedback structure is quite different to that used in the proof of Theorem 1, as discussed later.

Consider the error,

$$S_t^{s_0} - \widehat{S}_t^{s_0} = \Psi_t \cdots \Psi_2 \Psi_1(s_0) - \widehat{\Psi}_t \cdots \widehat{\Psi}_2 \widehat{\Psi}_1(s_0) \quad (4.19)$$

for any $s_0 \in \mathbb{S}$ and with $\mathbf{s}_0 = \mathbf{s}_0^{(L-1)}$ in this case. Let $\mathbf{s}_0^{(L-1)} = s_0$. We expand this error as a telescopic sum as,

$$\begin{aligned} S_t^{s_0} - \widehat{S}_t^{s_0} &= \sum_{\tau=1}^t \left(\widehat{\Psi}_t \cdots \widehat{\Psi}_{\tau+1} \Psi_\tau(\Psi_{\tau-1} \cdots (s_0)) \right. \\ & \quad \left. - \widehat{\Psi}_t \cdots \widehat{\Psi}_{\tau+1} \widehat{\Psi}_\tau(\Psi_{\tau-1} \cdots (s_0)) \right) \end{aligned} \quad (4.20)$$

This expansion formula is easy to check, e.g. try $t \in \{3, 4\}$. Note $S_{\tau-1}^{s_0} = \Psi_{\tau-1} \cdots (s_0)$, and we then have,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| S_t^{s_0} - \widehat{S}_t^{s_0} \right\|^2 \right]^{\frac{1}{2}} &= \\ & \leq \sum_{\tau=1}^t \mathbf{E}^\mu \left[\left\| \widehat{\Psi}_t \cdots \widehat{\Psi}_{\tau+1} \Psi_\tau(S_{\tau-1}^{s_0}) \right. \right. \\ & \quad \left. \left. - \widehat{\Psi}_t \cdots \widehat{\Psi}_{\tau+1} \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.21)$$

Assume now there exists real finite constants $C > 0$, $0 < \kappa < 1$ independent of $\tau \leq t \in \mathbb{N}$ such that,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| \widehat{\Psi}_t \cdots \widehat{\Psi}_\tau(\widehat{S}_{\tau-1}^{s_0}) - \widehat{\Psi}_t \cdots \widehat{\Psi}_\tau(\widehat{S}_{\tau-1}^s) \right\|^2 \right]^{\frac{1}{2}} \\ \leq C \kappa^{(t-\tau)} \mathbf{E}^\mu \left[\left\| \widehat{S}_{\tau-1}^{s_0} - \widehat{S}_{\tau-1}^s \right\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.22)$$

for any initial points $s_0, s \in \mathbb{S}$ and all $\tau, t \in \mathbb{N}$. That is, we assume the approximated filter inherits the hypothesised stability of the optimal filter (maybe with different constants, but the specifics of the constants won’t matter). We verify this assumption later. Applying this condition,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| S_t^{s_0} - \widehat{S}_t^{s_0} \right\|^2 \right]^{\frac{1}{2}} &= \\ & \leq \mathbf{E}^\mu \left[\left\| \Psi_t(S_{t-1}^{s_0}) - \widehat{\Psi}_t(S_{t-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \\ & \quad + C \sum_{\tau=1}^{t-1} \lambda^{(t-\tau-1)} \mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s_0}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.23)$$

we get a sum of one-step RNN-based approximation errors, with each term weighted by the stability factor of $\widehat{\Psi}_t$.

Let S' be some random variable with measure $\eta' \in \mathcal{P}(\mathbb{S})$. We consider $\mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s_0}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\| \right]$ for any $\tau \in$

$\{1, \dots, t\}$. Applying the triangle inequality twice we get,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s_0}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \\ \leq \mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s_0}) - \Psi_\tau(S_{\tau-1}^{s'_1}) \right\|^2 \right]^{\frac{1}{2}} \\ + \mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s'_1}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s'_1}) \right\|^2 \right]^{\frac{1}{2}} \\ + \mathbf{E}^\mu \left[\left\| \widehat{\Psi}_\tau(S_{\tau-1}^{s'_1}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.24)$$

From (4.13) and (4.22), with the latter to be verified, the first and third terms in the last inequality satisfy,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| \Psi_\tau(S_{\tau-1}^{s_0}) - \Psi_\tau(S_{\tau-1}^{s'_1}) \right\|^2 \right]^{\frac{1}{2}} \\ + \mathbf{E}^\mu \left[\left\| \widehat{\Psi}_\tau(S_{\tau-1}^{s'_1}) - \widehat{\Psi}_\tau(S_{\tau-1}^{s_0}) \right\|^2 \right]^{\frac{1}{2}} \\ \leq c \lambda^\tau \mathbf{E}^\mu \left[\|s_0 - s'_1\|^2 \right]^{\frac{1}{2}} \end{aligned} \quad (4.25)$$

for all $\tau \in \{1, \dots, t\}$ and some finite $c > 0$, $0 < \lambda < 1$. Let $C_0 := c \mathbf{E}^\mu [\|s_0 - s'_1\|^2]^{\frac{1}{2}} < \infty$.

It follows by a (non-trivial, but classical) result of Elton (1990) and an extension in Debaly and Truquet (2021), that given Assumption 3 and stationarity of the observation sequence $(Y_t)_{t \in \mathbb{N}}$, there exists a random S_∞ such that $(S_t^\eta)_{t \in \mathbb{N}}$ with $\eta := \text{Law}(S_\infty) \in \mathcal{P}(\mathbb{S})$ is stationary in \mathbb{S} .

With $\eta' = \eta$, note $S_\tau^{s'} \sim \eta, \forall \tau \in \mathbb{N}$ by stationarity. Then,

$$\begin{aligned} \mathbf{E}^\mu \left[\left\| \widehat{S}_t^{s_0} - S_t^{s_0} \right\|^2 \right]^{\frac{1}{2}} \leq \left(\frac{C + 1 - \kappa}{1 - \kappa} \right) \times \\ \left(C_0 + \mathbf{E}^\mu \left[\left\| \Psi(S, Y) - \widehat{\Psi}(S, Y) \right\|^2 \right]^{\frac{1}{2}} \right) \end{aligned} \quad (4.26)$$

where S has distribution η , and Y is distributed according to the invariant law of the stationary observation sequence.

Bounding $\mathbf{E}^\mu [\|\Psi(S, Y) - \widehat{\Psi}(S, Y)\|^2]$ is then achieved via classical universal approximation results for feedforward neural networks. In particular, applying Kidger and Lyons (2020, Theorem 4.16) we know a network exists such that,

$$\mathbf{E}^\mu \left[\left\| \Psi(S, Y) - \widehat{\Psi}(S, Y) \right\|^2 \right]^{\frac{1}{2}} \leq \varepsilon \quad (4.27)$$

for any $\varepsilon > 0$, with the added property that the Lipschitz constant of $\widehat{\Psi}(\cdot, y)$ is less than or equal to that of $\Psi(\cdot, y)$. The Lipschitz property follows from the main result in Eckstein (2020), see also Neumayer et al. (2022), when combined with Kidger and Lyons (2020, Proposition 4.9) in the proof of Kidger and Lyons (2020, Theorem 4.16).

Any $\varepsilon > 0$ that solves $\varepsilon \geq (\varepsilon + C_0) \frac{C+1-\kappa}{1-\kappa}$ is enough for the desired result (4.14).

It remains to establish the assumed (4.22). However, this follows directly from the fact that the Lipschitz constant of the chosen $\widehat{\Psi}(\cdot, y)$ is less than or equal to that of $\Psi(\cdot, y)$. In which case, (4.22) holds whenever (4.13) holds. \square

4.3 Discussion

Practical consequences: The results above have relevant practical implications. For example, Theorem 2 implies that RNN approximations to the optimal filter may be accurately applied on test sequences *indefinitely*, even when trained on finite-length data sequences (maybe only a handful of training time steps). This is desirable in online filtering applications. In particular, the approximation errors made at each step do not accumulate unbounded over time.

Note that Theorem 2 is divorced from any training procedure and is an idealised result. In practice, one is unlikely to know what a suitable sufficient statistic looks like and may only be interested in finding a small number of optimal conditional moments (e.g. maybe just the mean and covariance). A loss function used in practice for training a network might then only account for a small subset of the sufficient statistics required to define an optimal filtering recursion. This does not limit the theory. The recursion on the (wide) final activation layer may propagate a much higher-dimensional statistic defining the optimal filter than is carried forward through a linear output layer. In practice one may also consider a deeper feedforward network appended after the feedback layer, so as to compute other conditional functions. While these design and training ideas are not pertinent to the stated approximation capability results, they may be relevant in practical network design and in the design of loss functions and (e.g. hierarchical) training schemes.

The feedback in (4.2), (4.3), (4.4), used in the proof of Theorem 1, acts to memorise the observations, and the network otherwise seeks to approximate (non-recursive) solutions to (2.4), as in (2.3). Conversely, the time-uniform result in Theorem 2 is based on a network that directly approximates the optimal filtering recursion itself. This distinction offers some insight into the contrast between memorisation versus recursive state feedback in RNNs. Contrasting these two structures may influence network design and training in some applications; e.g. in models with long time-dependencies, or if the signal is not well understood, or perhaps to counter the effects of finite truncation in the number of sufficient statistics. Although the network (4.2), (4.3), (4.4) involves a memorisation over the entire finite interval of interest, in practice $\mathbf{E}^\mu [\rho(X_\tau) | Y_\tau, \dots, Y_1]$ may depend only weakly on observations in the distant past, e.g. as recognised and exploited in so-called fixed-lag smoothing methods, see Moore (1973). Thus, a smaller network, e.g. for (4.12), with a combination of recursion and memorisation may provide good longer-time performance.

Limitations: Unfortunately, (4.14) cannot be expected to hold, in general domains, under much weaker assumptions. For example, consider a linear-Gaussian signal/observation model that is controllable and observable, but with an unstable latent signal. The Kalman filter is stable in this case

and Assumption 3, i.e. (4.13), holds. However, in this case, the transition of X_t is not ergodic and $\mathbf{P}^\mu(Y_t \in \cdot)$ places most mass on sets of ever increasing distance from the origin. Thus, Assumption 2 does not hold, and its not possible to uniformly control the one-step approximation error as needed in the latter parts of the proof of Theorem 2.

In general, there are limited circumstances in which a finite-dimensional recursive filter of sufficient statistics exists. Even if such a filter exists, the square-integrability assumption on the sufficient statistic itself may be limiting. However, there may be natural finite-dimensional truncations of infinite-dimensional representations in the form (2.6), e.g. see Goodman (1975); Marcus (1979). For example, a Taylor-series truncation of the optimal filter, up to any finite order, can be approximated with this framework (with the most basic truncation being the so-called EKF, see Goodman (1975); Anderson and Moore (1979), based on a first-order Taylor expansion).

Since the filter is in practice learnt from data, and large neural networks are seemingly able to accurately capture complex high-dimensional maps, a RNN algorithmic framework offers distinct advantages in Bayesian filtering and time-series inference (e.g. over existing methods like model-based filters like the EKF/UKF, or Monte Carlo/particle methods). For example, a RNN may capture much higher-order finite truncations of some infinite optimal filtering statistic than is possible in other approaches, with less (test-time) computational burden. Conversely, Theorem 2 sheds some light on the theoretical limitations, viz. Assumption 2, of (at least naive versions of) this method in approximating recursive filter models; though contrast these with Theorem 1 on finite intervals.

The broad approach to prove Theorem 2 follows from results in the particle filtering community due to Del Moral and Guionnet (2001). Related work is focused on the transfer of stability from the optimal filter to the (particle) approximations; with deep technical methods aimed at proving and applying conditions like in Assumption 3, see, e.g., Heine and Crisan (2008); van Handel (2009); Whiteley (2013); Douc et al. (2014) and the stability literature referenced earlier. See also Crisan et al. (2020). The setup and proof of Theorem 2 takes a simpler view in terms of Assumption 3, so as to maximise the pedagogical value of these proof methods in problems involving RNNs, where they may not be known. This approach also offers some insight into the specific formulation and obstacles native to the generic neural approximation, e.g. as per Assumption 2. The details here contrast with other approximations: e.g. with particle approximations the one-step error may be controlled easier over longer horizons without Assumption 2. However, general theoretical results with particles may still be limited by comparable assumptions, as in van Handel (2009); Douc et al. (2014). Indeed, in Bishop and Del Moral (2020) it is proven that the basic bootstrap parti-

cle filter of Gordon et al. (1993) is incapable of tracking an unstable linear-Gaussian signal; in which setting Assumption 2 fails to hold, although Assumption 3 still holds.

Extensions: A trained RNN-based estimator will apply in practice under more complicated models than posited earlier; e.g., non-independent, non-Markov or time-varying.

Exponential ergodicity of the latent signal is in general settings likely sufficient for Theorem 2 to hold; i.e. this property implies asymptotic stationarity of the observations, and likely transfers to the filter so conditions like Assumption 3 hold. Next we study a special case in which Assumption 2 can be relaxed to just asymptotic stationarity.

4.4 Extension to Asymptotically Stationary Observations: Kalman Filtering Case Study

We examine a mild, but useful relaxation of Assumption 2; and illustrate where problems may arise when this assumption does not hold. Consider a linear-Gaussian model,

$$X_t^\mu = \mathbf{F}X_{t-1}^\mu + V_t \quad (4.28)$$

$$Y_t = \mathbf{H}X_t + W_t \quad (4.29)$$

where \mathbf{F} and \mathbf{H} are real matrices, and X_0 has Gaussian distribution $\mu \in \mathcal{P}(\mathbb{X})$ with mean \bar{X}_0 and covariance \mathbf{C}_0 . The random sequences $(V_t)_{t \in \mathbb{N}}$, $(W_t)_{t \in \mathbb{N}}$ are mutually independent zero-mean Gaussian with covariance matrices $\mathbf{Q} \geq 0$ and $\mathbf{R} > 0$, also independent of X_0 . If the model is detectable and stabilisable, see Anderson and Moore (1979), then Assumption 3 holds. We assume even more later.

We define the filter state $S_t^{(\bar{X}_0, \mathbf{C}_0)} := (\bar{X}_t, \mathbf{C}_t) \in \mathbb{S}$ via a slight abuse of notation. The form of the true Kalman filter for the model (4.28), (4.29) is given by,

$$\begin{aligned} \bar{X}_t &:= \mathbf{E}^\mu [X_t | Y_1, \dots, Y_t] \\ &= (\mathbf{F} - \mathbf{K}_t^* \mathbf{H} \mathbf{F}) \bar{X}_{t-1} + \mathbf{K}_t^* Y_t \end{aligned} \quad (4.30)$$

$$\begin{aligned} \mathbf{C}_t &:= \mathbf{Cov}^\mu [X_t | Y_1, \dots, Y_t] = \mathbf{Cov}^\mu [X_t - \bar{X}_t] \\ &= (\mathbf{I} - \mathbf{K}_t^* \mathbf{H}) (\mathbf{F} \mathbf{C}_{t-1} \mathbf{F}^\top + \mathbf{Q}) \end{aligned} \quad (4.31)$$

where $\mathbf{Cov}^\mu[\cdot]$ denotes the relevant covariance, and,

$$\begin{aligned} \mathbf{K}_t^* &= (\mathbf{F} \mathbf{C}_{t-1} \mathbf{F}^\top + \mathbf{Q}) \mathbf{H}^\top \times \\ &\quad [\mathbf{R} + \mathbf{H} (\mathbf{F} \mathbf{C}_{t-1} \mathbf{F}^\top + \mathbf{Q}) \mathbf{H}^\top]^{-1} \end{aligned} \quad (4.32)$$

Corollary 1. Consider a linear-Gaussian signal and observation model (4.28), (4.29). Assume the spectral radius of \mathbf{F} is strictly < 1 . Let $\hat{\Psi}^\theta(\mathbf{s}_{t-1}, Y_t)$ denote a RNN, as in (3.1), (3.2), (3.3), taking as input elements in the sequence $(Y_t)_{t \in \mathbb{N}}$. Then for any $\epsilon > 0$, there is a finite real parameter θ^* , and an initialisation vector \mathbf{s}_0 , such that,

$$\mathbf{E}^\mu \left[\left\| S_t^{(\bar{X}_0, \mathbf{C}_0)} - \hat{\Psi}^{\theta^*}(\mathbf{s}_{t-1}, Y_t) \right\|^2 \right]^{\frac{1}{2}} \leq \epsilon \quad (4.33)$$

for all $t \in \mathbb{N}$ and all $(\bar{X}_0, \mathbf{C}_0) \in \mathbb{S}$

The observation sequence is neither Markov nor stationary in general, and the hypotheses of the corollary here do not call for these conditions.

Proof. Stacking (4.30) and (4.29) we may write,

$$\begin{aligned} \begin{bmatrix} \bar{X}_t \\ X_t \end{bmatrix} &= \begin{bmatrix} \mathbf{F} - \mathbf{K}_t^* \mathbf{H} \mathbf{F} & \mathbf{K}_t^* \mathbf{H} \mathbf{F} \\ \mathbf{0} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \bar{X}_{t-1} \\ X_{t-1} \end{bmatrix} \\ &+ \begin{bmatrix} \mathbf{K}_t^* & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} W_t \\ V_t \end{bmatrix} \end{aligned} \quad (4.34)$$

The transition matrix for the given stacked system has spectral radius < 1 (its eigenvalues are the union of those of the diagonal blocks). We consider a RNN approximation of (4.34) and (4.31), treating this pair as the filter of interest. Assumption 3 holds for the pair (4.34), (4.31). The added noise process for the stacked system is stationary and thus Assumption 2 holds. Theorem 2 can now be applied. That is, there is a RNN that may approximate the system (4.34), (4.31) with (W_t, V_t) as an input. Suppose that we substitute Y_t for the input W_t , drop the input V_t , and set to zero the affine feedback map of \bar{X}_t . The result is a sub-network approximation of the desired (4.30), (4.31). By enforcing an accurate enough network approximation of (4.34), the result (4.33) holds for the sub-network approximation of the desired (4.30) when swapping the input from W_t to Y_t . This follows because,

$$\begin{aligned} \mathbf{E}^\mu[\|Y_t - W_t\|^2]^{\frac{1}{2}} &\leq \mathbf{E}^\mu[\|\mathbf{H}\mathbf{F}^t X_0 + \mathbf{H}\sum_{k=1}^{t-1} \mathbf{F}^k V_k\|^2]^{\frac{1}{2}} \\ &\leq \|\mathbf{H}\| \|\mathbf{F}^t\| (\text{trace}(\mathbf{C}_0) + \|\bar{X}_0\|^2)^{\frac{1}{2}} \\ &\quad + \frac{\|\mathbf{H}\| \text{trace}(\mathbf{Q})^{\frac{1}{2}}}{1 - \|\mathbf{F}\|} \end{aligned} \quad (4.35)$$

is just a finite constant (and $\|\mathbf{F}^t\|$ even goes to zero). \square

We remark that it is just the one model property (i.e. the spectral radius of \mathbf{F} being < 1) that ensures the observation sequence is asymptotically stationary, and that Assumption 3 is satisfied. This supports the more general notion that exponential ergodicity of the latent signal is sufficient for Theorem 2 to hold, as noted just before this subsection. Note Assumption 3 holds under just model detectability and stabilisability conditions (Anderson and Moore, 1979) (which can hold when the signal is unstable; but which are implied when the spectral radius of \mathbf{F} is < 1).

We have mildly relaxed the Assumption 2 that $(Y_t)_{t \in \mathbb{N}}$ is stationary used in Theorem 2. To see where things may go wrong, consider an identity approximation of a scalar observation Y_t given by $\sigma(Y_t + b) - b$ for a sufficiently large scalar bias $b > 0$. If Y_t is stationary, then there is a b such that $\mathbf{E}^\mu[\|Y_t - (\sigma(Y_t + b) - b)\|^2] \leq \varepsilon$ for any $\varepsilon > 0$. However, suppose X_t is unstable and Y_t moves to the negative; then there is a time $\tau = \tau(b) \in \mathbb{N}$ such that with exponentially large probability we have $\sigma(Y_t + b) - b = -b$ for all $t > \tau$. The approximation error for all $t > \tau$ grows on average and a time-uniform bound is impossible.

5 EXAMPLE

We consider a simple scalar example to illustrate the results, conditions, and some of the discussion points. Let,

$$\begin{aligned} X_t &= \alpha X_{t-1} + V_t, \\ Y_t &= X_t + \beta W_t \end{aligned} \quad (5.1)$$

where $\mathbb{X} = \mathbb{R}$, $\alpha \in \{0.98, 1.001\}$, $\beta \in \{1, 2\}$ and V_t, W_t are independent standard Gaussian white noises, independent of X_0 which is zero mean Gaussian with variance 25.

The Kalman filter $S_t^{(\bar{X}_0, \mathbf{C}_0)} := (\bar{X}_t, \mathbf{C}_t)$ is optimal and Assumption 3 holds in all cases considered (by satisfaction of the classical detectability and stabilisability conditions). If $|\alpha| < 1$, then X_t is ergodic and Assumption 2 holds. Here, $|\alpha| < 1$ is an explicit provable threshold on the satisfaction of Assumption 2. We also have the ground-truth comparison in the optimal Kalman filter. We can thus study the requisite condition for time-uniform approximation (e.g. moving from $\alpha = 0.98$ to just $\alpha = 1.001$).

We train a RNN based on (3.6) with $\rho(x) = x$, i.e. neglecting the variance. We use N_{train} samples over horizons of length T_{train} . This training loss function is consistent with prior discussion noting that in practice one is unlikely to know the nature of sufficient statistics, nor desire them all as outputs, nor be willing to train on loss functions with all of them. We also consider a basic particle filter implementation (Gordon et al., 1993) with 1000 particles.

We compute $(\frac{1}{N_{\text{test}}} \sum_{n=1}^{N_{\text{test}}} (\bar{X}_t^{(n)} - \hat{X}_t^{(n)})^2)^{\frac{1}{2}}$ where $\hat{X}_t^{(n)}$ is the n -th mean state estimate of the particle or RNN approximation from $N_{\text{test}} = 1000$ independent test examples. The test horizon is $T_{\text{test}} = 2000$ in all cases.

In Fig 1 ($\alpha = 0.98$, $\beta \in \{1, 2\}$) and Fig 2 ($\alpha = 1.001$, $\beta = 2$) we plot the errors for different parameters. For each case we consider a RNN in the form (4.15), (4.16), (4.17) with $L = 3$ and 7 RELUs on both hidden layers.

In each case in Fig 1 we train with just $T_{\text{train}} = 20$, and $N_{\text{train}} = 5000$. With $\alpha = 0.98$, Corollary 1 holds, and the RNN-based method outperforms the particle filter (which does not handle accurate observations $\beta = 1$ well here).

In Fig 2 we consider $T_{\text{train}} = 20$, $N_{\text{train}} = 5000$ (as before), and also $T_{\text{train}} = 200$, $T_{\text{train}} = 1000$, $T_{\text{train}} = 2000$ (in the latter two cases we reduce $N_{\text{train}} = 1000$).

These examples reflect on a number of discussion points: 1) training on an output/cost function of interest, but of lower dimension than the sufficient statistic required to define the optimal filter; 2) when the conditions of Theorem 2/Corollary 1 are met, we may train on finite-length (often very short, maybe one or a handful of steps, depending on the ergodicity) but apply the filter indefinitely at test time with no unbounded accumulated growth of error; 3) differing capabilities and particulars of differing approximation




Figure 1: Root mean square errors.




Figure 2: Root mean square errors (log vertical scale).

schemes (e.g. particle vs neural in this case); 4) if either assumption is unmet, we may learn networks that work well on the length of the training data, e.g. as in Theorem 1, but eventually the error may start to accumulate.

6 CONCLUSION

We consider a generic recurrent neural network framework that approximates directly a recursive mapping from observational inputs to some desired Bayesian filter statistics. The focus of this article is the approximation capability of this framework. We provide approximation error bounds for filtering in general non-compact domains. The main result of this work is a couple of conditions on the underlying observation sequence and on the optimal filter that when

satisfied allow one to approximate the optimal Bayesian filter to any desired accuracy over an indefinite or infinite time horizon. When applicable, this strong, time-uniform approximation result ensures good long-time filtering performance in practice. We discuss and illustrate a number of practical concerns and implications of these results; and we contrast the fixed horizon and time-uniform results with each other. With respect to the time-uniform universal approximation result, we explore the mechanisms by which the required conditions manifest in practice, their necessity, and we discuss their appearance in similar results based on different filtering approximation schemes.

References

- Alessandri, A., Baglietto, M., Parisini, T., and Zoppoli, R. (1999). A neural state estimator with bounded errors for nonlinear systems. *IEEE Transactions on Automatic Control*, 44(11):2028–2042.
- Anderson, B. D. O. and Moore, J. B. (1979). *Optimal Filtering*. Dover Publications.
- Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342.
- Atar, R. and Zeitouni, O. (1997). Lyapunov Exponents for Finite State Nonlinear Filtering. *SIAM Journal on Control and Optimization*, 35(1):36–55.
- Bain, A. and Crisan, D. (2008). *Fundamentals of Stochastic Filtering*. Stochastic Modelling and Applied Probability. Springer-Verlag GmbH.
- Becker, P., Pandya, H., Gebhardt, G., Zhao, C., Taylor, C. J., and Neumann, G. (2019). Recurrent Kalman networks: Factorized inference in high-dimensional deep feature spaces. In *International Conference on Machine Learning*, pages 544–552. PMLR.
- Bishop, A. N. and Del Moral, P. (2017). On the Stability of Kalman-Bucy Diffusion Processes. *SIAM Journal on Control and Optimization*, 55(6):4015–4047.
- Bishop, A. N. and Del Moral, P. (2020). On the mathematical theory of ensemble (linear-Gaussian) Kalman-Bucy filtering. *arXiv e-print: 2006.08843*.
- Budhiraja, A. and Ocone, D. (1999). Exponential stability in discrete-time filtering for non-ergodic signals. *Stochastic Processes and their Applications*, 82(2):245–257.
- Chigansky, P. and Liptser, R. (2004). Stability of nonlinear filters in nonmixing case. *The Annals of Applied Probability*, 14(4).
- Crisan, D. and Heine, K. (2008). Stability of the discrete time filter in terms of the tails of noise distributions. *Journal of the London Mathematical Society*, 78(2):441–458.
- Crisan, D., López-Yela, A., and Miguez, J. (2020). Stable Approximation Schemes for Optimal Filters. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):483–509.
- Daum, F. E. (1986). Exact finite-dimensional nonlinear filters. *IEEE Transactions on Automatic Control*, 31(7):616–622.
- Debaly, Z. M. and Truquet, L. (2021). Iterations of dependent random maps and exogeneity in nonlinear dynamics. *Econometric Theory*, 37(6):1135–1172.
- Del Moral, P. and Guionnet, A. (2001). On the stability of interacting processes with applications to filtering and genetic algorithms. *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, 37(2):155–194.
- Dickinson, B. W. and Sontag, E. D. (1985). Dynamic realizations of sufficient sequences. *IEEE Transactions on Information Theory*, 31(5):670–676.
- Doerr, A., Daniel, C., Schiegg, M., Duy, N.-T., Schaal, S., Toussaint, M., and Sebastian, T. (2018). Probabilistic recurrent state-space models. In *International Conference on Machine Learning*, pages 1280–1289. PMLR.
- Douc, R., Moulines, E., and Olsson, J. (2014). Long-term stability of sequential Monte Carlo methods under verifiable conditions. *The Annals of Applied Probability*, 24(5).
- Eckstein, S. (2020). Lipschitz neural networks are dense in the set of all Lipschitz functions. *arXiv e-print:2009.13881*.
- Elton, J. H. (1990). A multiplicative ergodic theorem for Lipschitz maps. *Stochastic Processes and their Applications*, 34(1):39–47.
- Evensen, G. (2009). *Data Assimilation: The Ensemble Kalman Filter*. Springer, second edition.
- Fox, E., Sudderth, E., Jordan, M., and Willsky, A. (2008). Nonparametric Bayesian learning of switching linear dynamical systems. In *Advances in Neural Information Processing Systems*, volume 21.
- Frerix, T., Kochkov, D., Smith, J., Cremers, D., Brenner, M., and Hoyer, S. (2021). Variational data assimilation with a learned inverse observation operator. In *International Conference on Machine Learning*, pages 3449–3458. PMLR.
- Frigola, R., Chen, Y., and Rasmussen, C. E. (2014). Variational Gaussian Process State-Space Models. In *Advances in Neural Information Processing Systems*, volume 27.
- Gerber, M. and Whiteley, N. (2017). Stability with respect to initial conditions in V-norm for nonlinear filters with ergodic observations. *Journal of Applied Probability*, 54(1):118–133.
- Ghahramani, Z. and Hinton, G. E. (2000). Variational learning for switching state-space models. *Neural Computation*, 12(4):831–864.
- Ghahramani, Z. and Jordan, M. (1995). Factorial hidden Markov models. In *Advances in Neural Information Processing Systems*, volume 8.
- Goodman, I. R. (1975). Some Results on the Additive Nonlinear Discrete Filtering Problem. In *Proceedings of the 6th Symposium Nonlinear Estimation and Its Applications*, pages 100–103.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F: Radar and Signal Processing*, 140(2):107–113.

- Gu, S. S., Ghahramani, Z., and Turner, R. E. (2015). Neural adaptive sequential Monte Carlo. In *Advances in Neural Information Processing Systems*, volume 28.
- Haarnoja, T., Ajay, A., Levine, S., and Abbeel, P. (2016). Backprop KF: Learning discriminative deterministic state estimators. In *Advances in Neural Information Processing Systems*, volume 29.
- Haykin, S., Yee, P., and Derbez, E. (1997). Optimum nonlinear filtering. *IEEE Transactions on Signal Processing*, 45(11):2774–2786.
- Heine, K. and Crisan, D. (2008). Uniform approximations of discrete-time filters. *Advances in Applied Probability*, 40(4):979–1001.
- Ialongo, A. D., Van Der Wilk, M., Hensman, J., and Rasmussen, C. E. (2019). Overcoming Mean-Field Approximations in Recurrent Gaussian Process Models. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2931–2940. PMLR.
- Jasra, A., Singh, S. S., Martin, J. S., and McCoy, E. (2010). Filtering via approximate Bayesian computation. *Statistics and Computing*, 22(6):1223–1237.
- Julier, S. and Uhlmann, J. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45.
- Karl, M., Soelch, M., Bayer, J., and van der Smagt, P. (2016). Deep Variational Bayes Filters: Unsupervised Learning of State Space Models from Raw Data. In *International Conference on Learning Representations*.
- Kidger, P. and Lyons, T. (2020). Universal Approximation with Deep Narrow Networks. In *Proceedings of Thirty Third Conference on Learning Theory*, pages 2306–2327. PMLR.
- Kleptsyna, M. L. and Veretennikov, A. Y. (2007). On discrete time ergodic filters with wrong initial data. *Probability Theory and Related Fields*, 141(3-4):411–444.
- Ko, J. and Fox, D. (2009). GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models. *Autonomous Robots*, 27(1):75–90.
- Krishnan, R. G., Shalit, U., and Sontag, D. (2017). Structured Inference Networks for Nonlinear State Space Models. In *Proceedings of the Thirty First AAAI Conference on Artificial Intelligence*, pages 2101–2109.
- Levine, J. and Pignie, G. (1986). Exact finite dimensional filters for a class of nonlinear discrete-time systems. *Stochastics*, 18(2):97–132.
- Lo, J. T.-H. (1994). Synthetic approach to optimal filtering. *IEEE Transactions on Neural Networks*, 5(5):803–811.
- Marcus, S. (1979). Optimal nonlinear estimation for a class of discrete-time stochastic systems. *IEEE Transactions on Automatic Control*, 24(2):297–302.
- Martin, G. M., McCabe, B. P. M., Frazier, D. T., Manesoonthorn, W., and Robert, C. P. (2019). Auxiliary Likelihood-Based Approximate Bayesian Computation in State Space Models. *Journal of Computational and Graphical Statistics*, 28(3):508–522.
- Moore, J. B. (1973). Discrete-time fixed-lag smoothing algorithms. *Automatica*, 9(2):163–173.
- Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT press.
- Neumayer, S., Goujon, A., Bohra, P., and Unser, M. (2022). Approximation of Lipschitz Functions using Deep Spline Neural Networks. *arXiv e-print:2204.06233*.
- Nickisch, H., Solin, A., and Grigorevskiy, A. (2018). State space Gaussian processes with non-Gaussian likelihood. In *International Conference on Machine Learning*, pages 3789–3798. PMLR.
- Oudjane, N. and Rubenthaler, S. (2005). Stability and Uniform Particle Approximation of Nonlinear Filters in Case of Non Ergodic Signals. *Stochastic Analysis and Applications*, 23(3):421–448.
- Parisini, T., Alessandri, A., Maggiore, M., and Zoppoli, R. (1997). On convergence of neural approximate nonlinear state estimators. In *Proceedings of the 1997 American Control Conference*, pages 1819–1822. IEEE.
- Parisini, T. and Zoppoli, R. (1994). Neural networks for nonlinear state estimation. *International Journal of Robust and Nonlinear Control*, 4(2):231–248.
- Parlos, A. G., Menon, S. K., and Atiya, A. (2001). An algorithmic approach to adaptive state filtering using recurrent neural networks. *IEEE Transactions on Neural Networks*, 12(6):1411–1432.
- Pascanu, R., Gulcehre, C., Cho, K., and Bengio, Y. (2014). How to construct deep recurrent neural networks. In *Proceedings of the 2nd International Conference on Learning Representations*.
- Pitt, M. K. and Shephard, N. (1999). Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*, 94(446):590–599.
- Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. (2018). Deep State Space Models for Time Series Forecasting. In *Advances in Neural Information Processing Systems*, volume 31.
- Rudenko, E. A. (2010). Optimal discrete nonlinear filters of the objects’s order and their Gaussian approximations. *Automation and Remote Control*, 71(2):320–338.
- Sawitzki, G. (1981). Finite dimensional filter systems in discrete time. *Stochastics*, 5(1-2):107–114.

- Segall, A. (1976). Recursive estimation from discrete-time point processes. *IEEE Transactions on Information Theory*, 22(4):422–431.
- Stettner, L. (1989). On invariant measures of filtering processes. In Christopheit, N., Helmes, K., and Kohlmann, M., editors, *Stochastic Differential Systems*, volume 126 of *Lecture Notes in Control and Information Sciences*, pages 279–292. Springer-Verlag.
- Tong, X. T. and van Handel, R. (2012). Ergodicity and stability of the conditional distributions of nondegenerate Markov chains. *The Annals of Applied Probability*, 22(4).
- V. T., S. E. and Shin, Y. C. (1994). Radial basis function neural network for approximation and estimation of non-linear stochastic dynamic systems. *IEEE Transactions on Neural Networks*, 5(4):594–603.
- Van Der Merwe, R., Doucet, A., De Freitas, N., and Wan, E. (2000). The unscented particle filter. In *Advances in Neural Information Processing Systems*, volume 13.
- van Handel, R. (2008). Discrete time nonlinear filters with informative observations are stable. *Electronic Communications in Probability*, 13:562–575.
- van Handel, R. (2009). Uniform time average consistency of Monte Carlo particle filters. *Stochastic Processes and their Applications*, 119(11):3835–3861.
- Whiteley, N. (2013). Stability properties of some particle filters. *The Annals of Applied Probability*, 23(6).