



HAL
open science

Last layer state space model for representation learning and uncertainty quantification

Max Cohen, Maurice Charbit, Sylvain Le Corff

► **To cite this version:**

Max Cohen, Maurice Charbit, Sylvain Le Corff. Last layer state space model for representation learning and uncertainty quantification. IEEE Workshop Statistical Signal Processing, 2023, 10.48550/arXiv.2307.01566 . hal-04146454

HAL Id: hal-04146454

<https://hal.science/hal-04146454>

Submitted on 3 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Last layer state space model for representation learning and uncertainty quantification

Max Cohen
Télécom SudParis, CITI, TIPIC
Institut Polytechnique de Paris
maxjcohen@proton.me

Maurice Charbit
Accentra, Boulogne-Billancourt
maurice.charbit@accentra.fr

Sylvain Le Corff
Télécom SudParis, CITI, TIPIC
Institut Polytechnique de Paris

Abstract—As sequential neural architectures become deeper and more complex, uncertainty estimation is more and more challenging. Efforts in quantifying uncertainty often rely on specific training procedures, and bear additional computational costs due to the dimensionality of such models. In this paper, we propose to decompose a classification or regression task in two steps: a representation learning stage to learn low-dimensional states, and a state space model for uncertainty estimation. This approach allows to separate representation learning and design of generative models. We demonstrate how predictive distributions can be estimated on top of an existing and trained neural network, by adding a state space-based last layer whose parameters are estimated with Sequential Monte Carlo methods. We apply our proposed methodology to the hourly estimation of Electricity Transformer Oil temperature, a publicly benchmarked dataset. Our model accounts for the noisy data structure, due to unknown or unavailable variables, and is able to provide confidence intervals on predictions.

Index Terms—Recurrent neural networks, Representation learning, Uncertainty quantification, Sequential Monte Carlo.

I. INTRODUCTION

Recurrent Neural Networks (RNN) were first introduced as an efficient and convenient architecture to address short time dependencies problems. They have been consistently improved to develop longer term memory, and optimize their implementations [1], [2]. Current deep learning frameworks allow stacking arbitrary high number of recurrent layers, whose parameters are estimated by gradient descent through automated differentiation procedures, as shown in [3]. However, many critical applications, such as medical diagnosis or drug design discovery, require not only accurate predictions, but a good estimate of their uncertainty ([4], [5]). Fostering the dissemination of deep learning-based algorithms to such fields requires to design new approaches for uncertainty quantification.

Bayesian statistics are able to approximate the distributions of future observations and to provide uncertainty estimation [6]. Several architectures inspired by Variational Inference (VI, see [7]) emerged by considering latent states as random variables and approximating their posterior distribution. The authors of [8] built on a traditional recurrent architecture by modelling temporal dependencies between these latent random states. Results presented in [9] yield improved performances when considering local gradient information for computing the posterior. In [10], a prior model based on a Markov chain is

estimated in the latent space of an Auto Encoder in order to compute uncertainty estimation on the observation.

Sequential Monte Carlo (SMC) methods have also been successfully applied to Recurrent Neural Networks. Instead of computing a single latent vector at each time step, a set of particles representing the distribution of the latent space are propagated, and associated with importance weights. In [11], [12], the authors were able to model complex distributions on dependant data. We turn to [13] for an example using more complex neural architectures, such as the Transformer.

In [14], the authors considered weights as random variables and proposed approximations of their posterior distributions allowing more robust predictions. Such Bayesian neural networks have been proposed and studied in a variety of works, see for instance [15], [16]. However, these methods are computationally intensive for high dimensional models and we do not have statistical guarantees on their ability to capture the target posterior distribution, see [17].

Monte Carlo Dropout (MC Dropout) methods offer to capture uncertainty by leveraging Dropout during both training and evaluation tasks, producing variable predictions from a single trained recurrent model, see [18]. In the recent years, MC Dropout methods have been applied in many industrial fields, such as flight delay prediction [19] or molecular simulations [20]. Alternatively, ensemble methods consist in training distinct networks to obtain a combined prediction, as shown in [21]. However, these frequentist approaches fail to guarantee proper calibration of the model, as highlighted by [22], and suffer various limitations, see [23].

In an effort to provide an alternative strategy with limited computation overhead, [24] suggests splitting representation learning and uncertainty estimation to solve classification problems for independent data. A deep classifier is first trained to obtain task dependent representations of the data, on which ensemble models are fitted to approximate the distribution of the observations. Their experiments indicate that performing uncertainty estimation on the last-layer of the model outperforms baseline networks and is an appealing trade-off between computational cost and uncertainty quantification. However, as this method is restricted to independent data, it cannot be directly applied to time series.

Inspired by [24], we propose a last layer approach to split uncertainty quantification from representation learning,

in the context of dependent data. This new method combines high expressivity, quality uncertainty estimations and ease of training. Our main contributions are as follows.

- We propose a decoupled architecture composed of an arbitrary sequential model and a state space model layer.
- This last layer allows to introduce complex predictive distributions for the observations. Its parameters are estimated through approximate samplings using Sequential Monte Carlo methods, as the likelihood of the observations is not available explicitly in such a setting.
- Our methodology allows for arbitrary deep architectures, and does not suffer the overconfidence of frequentist approaches.

II. LAST LAYER DECOUPLING

Estimating the parameters of potentially high-dimensional models with unobserved (i.e. noisy) layers is a challenging task. We therefore propose to first train an input model following traditional deep learning approaches, then use Monte Carlo methods in a lower dimensional state space to account for uncertainty, with tractable and computationally efficient simulation-based methods. The two-stage training algorithm is presented in Algorithm 1, and the architecture of the model is described in Figure 1

In the following, for any sequence (a_m, \dots, a_n) with $n \geq m$, we use the short-hand notation $a_{m:n} = (a_m, \dots, a_n)$. Let $T \geq 1$ be a given time horizon. We consider a regression task with observations $Y_{1:T}$ associated with inputs $U_{1:T}$.

A. Representation learning

In this paper, we consider an arbitrary multi-layer neural network h_φ with unknown parameters φ , responsible for extracting high level features from the input time series:

$$\tilde{U}_{1:T} = h_\varphi(U_{1:T}), \quad \text{input model.}$$

We produce an estimate $\hat{\varphi}$ during the first training stage, by introducing an auxiliary function κ_ψ to model the observations as follows: for all $1 \leq k \leq T$, $Y_k = \kappa_\psi(Y_{k-1}, \tilde{U}_k) + \epsilon_k$ and $Y_0 = \kappa_\psi(\tilde{U}_0) + \epsilon_0$, where $(\epsilon_k)_{k \leq 0}$ are independent centered Gaussian random variables with unknown variance. The input model is trained on a simple deterministic regression task, by performing gradient descent on the mean squared error, leading to a first estimate of φ and ψ . We keep the estimated parameters $\hat{\varphi}$ while the auxiliary function κ_ψ and its parameters, only designed to model the observations, are discarded.

B. State space model

The next step is to define a state space model taking as input the previously extracted features $\tilde{U}_{1:T}$. Let $X_{1:T}$ a sequence of stochastic hidden states computed recursively and Y_t their associated predictions. For all $k \geq 1$, the model is defined as:

$$\begin{aligned} X_k &= g_\theta(X_{k-1}, \tilde{U}_k) + \eta_k, & \text{state model,} \\ Y_k &= f_\theta(X_k) + \epsilon_k, & \text{observation model,} \end{aligned}$$

where θ are the unknown real-valued parameters of the network (weights and biases) and f_θ and g_θ are nonlinear parametric functions. We chose $(\eta_k)_{k \geq 1}$ and $(\epsilon_k)_{k \geq 1}$ as two sequences of independent centered Gaussian random variables with covariance matrices Σ_x and Σ_y , although any distribution can be substituted.

This decoupled approach aims at reducing the number of parameters in θ , compared to φ , in order to estimate them using Sequential Monte Carlo methods. In the next section, we describe this second training procedure for the last layer only, by keeping $\hat{\varphi}$ fixed.




Fig. 1. Our architecture combining a generic input model with a state space model on the last layer.

III. SEQUENTIAL MONTE CARLO LAYER

In this section, we detail how to estimate the parameters θ , Σ_x and Σ_y in the model introduced in Section II-B, from a record of observations $Y_{1:T}$. This is challenging because the likelihood of the observations is not available explicitly, as it would require integrating over the hidden states $X_{1:T}$. Consequently, the score function is intractable. We propose to optimize a Monte Carlo estimator of this score function, using Fisher's identity [25]:

$$\nabla_\theta \log p_\theta(Y_{1:T}) = \mathbb{E}_\theta [\nabla_\theta \log p_\theta(X_{1:T}, Y_{1:T}) | Y_{1:T}], \quad (1)$$

where \mathbb{E}_θ designs the expectation under the model parameterized by θ (the dependency on the input $U_{1:T}$ is kept implicit here for better clarity). In the following paragraphs, we denote by $\Psi_{\mu, \Sigma}$ the Gaussian probability density function with mean vector μ and covariance matrix Σ .

A. Particle filter

The conditional distribution of $X_{1:T}$ given $Y_{1:T}$ is not available explicitly for a nonlinear state space model, but it can be approximated using a family of N particles $(\xi_{1:T}^\ell)_{\ell=1}^N$ associated with importance weights $(\omega_{1:T}^\ell)_{\ell=1}^N$. At $k = 0$, $(\xi_0^\ell)_{\ell=1}^N$ are sampled independently from $\rho_0 = \Psi_{0, \Sigma_x}$, and each particle ξ_0^ℓ is associated with the standard importance sampling weight $\omega_0^\ell \propto \Psi_{Y_0, \Sigma_y}(f_\theta(\xi_0^\ell))$. Then, for $k \geq 1$, using $\{(\xi_{k-1}^\ell, \omega_{k-1}^\ell)\}_{\ell=1}^N$, we sample pairs $\{(I_k^\ell, \xi_k^\ell)\}_{\ell=1}^N$ of indices and particles from the instrumental distribution:

$$\pi_k(\ell, x) \propto \omega_{k-1}^\ell p_k(\xi_{k-1}^\ell, \tilde{U}_k, x).$$

In this application we use for $p_k(\xi_{k-1}^\ell, \tilde{U}_k, \cdot)$ the prior kernel $\Psi_{g_\theta(\xi_{k-1}^\ell, \tilde{U}_k), \Sigma_x}$. For $\ell \in \{1, \dots, N\}$, ξ_k^ℓ is associated with the

importance weight $\omega_k^\ell \propto \Psi_{Y_k, \Sigma_y}(f_\theta(\xi_k^\ell))$. Such a particle filter with multinomial resampling is referred to as the bootstrap algorithm, see [26]. It has been extended and analyzed in many directions in the past decades, see [27]–[29]. In other lines of works, the adaptive tuning of the Monte Carlo effort has been analyzed in order to adapt the number of particles on-the-fly, see [30], [31].

B. Particle smoother and online estimation

Our framework allows the use of any particle smoother to estimate (1). In this paper, we first describe the Path-space smoother [32] for its simplicity, in order to illustrate our approach. In practice, it often leads to particle path degeneracy [33], which can be mitigated by substituting a more complex smoother such as the Forward Filtering Backward Smoothing [34] or the Forward Filtering Backward Simulation algorithm [35]. Additionally, because estimating (1) amounts to computing a smoothed expectation of an additive functional, we can also use very efficient forward-only SMC smoothers such as the PaRIS algorithm and its pseudo-marginal extensions [36], [37]. With $\xi_{1:T}^i$ the ancestral line of ξ_T^i , the score function (1) can be estimated as follows using automated differentiation:

$$\widehat{S}_\theta^N(Y_{1:T}) = \sum_{\ell=1}^N \omega_T^\ell \nabla_\theta \log p_\theta(\xi_{1:T}^\ell, Y_{1:T}),$$

where p_θ is the joint probability density function of $(X_{1:T}, Y_{1:T})$ for the model described in Section II-B.

The degeneracy relative to the smoothing problem can be overcome using *backward sampling*. It is specifically designed for additive functionals so it is well suited to our setting (1) since $\nabla_\theta \log p_\theta(x_{1:T}, y_{1:T}) = \sum_{t=1}^T \nabla_\theta \log m_\theta(x_{t-1}, \tilde{u}_t; x_t) r_\theta(x_t, y_t)$, where $m_\theta(x_{t-1}, \tilde{u}_t; \cdot)$ is the transition density of the state model and $r_\theta(x_t, \cdot)$ is the density of the conditional distribution of y_t given x_t and by convention $m_\theta(x_0, \tilde{u}_1; \cdot) = \rho_0(\cdot)$.

The Monte Carlo estimator of the score function can be obtained online by setting, $\widehat{S}_\theta^N(y_{1:T}) = \sum_{i=1}^N \omega_T^i \tau_T^i$, where the statistics $\{\tau_s^i\}_{i=1}^N$ satisfy the recursion $\tau_{s+1}^i = \tau_s^{i_{s+1}} + \tilde{h}_s(\xi_s^{i_{s+1}}, \xi_{s+1}^i)$, where $\tilde{h}_s(x_s, x_{s+1}) = \nabla_\theta \log m_\theta(x_s, \tilde{u}_{s+1}; x_{s+1}) r_\theta(x_{s+1}, y_{s+1})$. Following [36]–[39], the degeneracy of the path-space smoother can be overcome by performing an online *PaRis* update of the statistics τ_{s+1}^i , $1 \leq i \leq N$, using the backward kernel of the hidden Markov chain.

An appealing application of the last layer approach is recursive maximum likelihood estimation, i.e., where new observations are used only once to update the estimator of the unknown parameter θ . In [24], the authors used in particular Stochastic Gradient Descent (SGD) and Stochastic Gradient Langevin Dynamics to update the estimation of θ and perform uncertainty quantification. In state space models, recursive maximum likelihood estimation produces a sequence $\{\theta_k\}_{k \geq 0}$ of parameter estimates writing, for each new observation Y_k , $k \geq 1$,

$$\theta_k = \theta_{k-1} + \gamma_k \nabla_\theta \ell_\theta(Y_k | Y_{0:k-1}),$$

where $\ell_\theta(Y_k | Y_{0:k-1})$ is the loglikelihood for the new observation given all the past, and $\{\gamma_k\}_{k \geq 1}$ are positive step sizes such that $\sum_{k \geq 1} \gamma_k = \infty$ and $\sum_{k \geq 1} \gamma_k^2 < \infty$. The practical implementation of such an algorithm, where $\nabla_\theta \ell_\theta(Y_k | Y_{0:k-1})$ is approximated using the weighted samples $\{(\xi_k^\ell, \omega_k^\ell)\}_{\ell=1}^N$ can be found for instance in [37]. The PaRIS algorithm proposed in [36] allows to use the weighted samples $\{(\xi_k^\ell, \omega_k^\ell)\}_{\ell=1}^N$ and the statistics $\{\tau_k^{\ell, j}\}_{\ell=1}^N$ on-the-fly to approximate $\nabla_\theta \ell_\theta(Y_k | Y_{0:k-1})$.

Although this algorithm is very efficient to update parameters recursively, it is computationally intensive and therefore fits particularly well our last layer approach as it would be intractable for very high dimensional latent states.

Algorithm 1: Two-stage learning

```

 $\hat{\varphi} \leftarrow$  Train the input model  $h_\varphi$ ;
 $\tilde{U}_{1:T} \leftarrow h_{\hat{\varphi}}(U_{1:T})$ ;
Initialize parameter estimate  $\hat{\theta}_0$ ;
for  $p \leftarrow 1$  to MaxIt do
   $\xi_0^\ell \sim \rho_0$  and  $\omega_0^\ell \propto \Psi_{Y_0, \Sigma_y}(f_{\hat{\theta}_{p-1}}(\xi_0^\ell))$ ;
  for  $k \leftarrow 1$  to  $T$  do
    for  $j \leftarrow 1$  to  $N$  do
       $I_k^j \sim \mathbb{P}(I_k^j = m) = \omega_{k-1}^m$ ;
       $\xi_k^j \sim p_k(\xi_{k-1}^{I_k^j}, \tilde{U}_k, \cdot)$ ;
       $\omega_k^j \propto \Psi_{Y_k, \Sigma_y}(f_{\hat{\theta}_{p-1}}(\xi_k^j))$ ;
      Set  $\xi_{0:k}^j = (\xi_{0:k-1}^{I_k^j}, \xi_k^j)$ .
    Update the parameter estimate using gradient
    descent with estimated gradient  $\widehat{S}_{\hat{\theta}_{p-1}}^N(Y_{1:T})$ .

```

IV. EXPERIMENTS

A. Data and model

We benchmarked our approach on the public Electricity Transformer Temperature (ETT) Dataset, designed in [40] to forecast Oil temperature based on hourly power load records (ETTh1 subset).

The Input model is a $L = 3$ layered GRU model, as defined in the deep learning framework PyTorch¹: for all $1 \leq \ell \leq L$ and all $1 \leq k \leq T$,

$$\begin{aligned} r_k^\ell &= \sigma(W_{ir} U_k^{\ell-1} + b_{ir} + W_{hr} U_{k-1}^\ell + b_{hr}), \\ z_k^\ell &= \sigma(W_{iz} U_k^{\ell-1} + b_{iz} + W_{hz} U_{k-1}^\ell + b_{hz}), \\ n_k^\ell &= \tanh(W_{in} U_k^{\ell-1} + b_{in} + r_k^\ell (W_{hn} U_{k-1}^\ell + b_{hn})), \\ \tilde{U}_k^\ell &= (1 - z_k^\ell) n_k^\ell + z_k^\ell U_{k-1}^\ell, \end{aligned}$$

where $\varphi = \{(W_{is}, b_{is}, W_{hs}, b_{hs}), s \in \{r, z, n\}\}$ are unknown parameters, and $\sigma : x \mapsto 1/(1 + e^{-x})$ is the sigmoid function. The first layer of the network is assimilated to the input vectors, $\tilde{U}_t^0 \equiv U_k$ and $\tilde{U}_0^\ell \equiv 0$. The input dimension $d_{in} = 6$ corresponds to the number of power load records of the dataset, we set the output dimension to 6. In order to estimate the parameters φ , we introduce an

¹<https://pytorch.org/docs/stable/generated/torch.nn.GRU.html>

auxiliary GRU layer responsible for computing oil temperature predictions. During the training, we minimize the cost function $\mathcal{L}_{\text{input}}(\varphi) = \sum_{i=1}^{N_{\text{sample}}} \|\text{model}_{\varphi}(U_{1:T}^i) - Y_{1:T}^i\|^2$ between each sample of the dataset and the associated prediction obtain with this deterministic model.

The State Space model is implemented using PyTorch implementations of RNN and Linear layers. We chose the following form for f_{θ} and g_{θ} :

$$\begin{aligned} g_{\theta} : X_{k-1}, \tilde{U}_k &\mapsto \tanh(W_{gx}X_{k-1} + b_{gx} + W_{gu}\tilde{U}_k + b_{gu}), \\ f_{\theta} : X_k &\mapsto \sigma(W_f X_k + b_f), \end{aligned}$$

where $\theta = \{W_{gx}, b_{gx}, W_{gu}, b_{gu}, W_f, b_f, \Sigma_x, \Sigma_y\}$ are unknown parameters. All following experiments are conducted with $N = 100$ particles and a batch size of 32, using the Adam optimizer introduced in [41]. The learning rate was chosen using a simple grid search. We train models for a maximum of 50 epochs, and employ early stopping to prevent overfit.

B. Evaluations

In this section, we illustrate the ability of our model to capture the distribution of future observations, by evaluating the benchmarked models using the following protocol. We draw 48 hours long samples $(u_{1:48}, y_{1:48})$ from the validation dataset, composed of a 24 hour long lookback window $(u_{1:24}, y_{1:24})$, containing historic commands and observations, and a predictions window where only future commands are available $(u_{25:48})$. Each model produces $N = 100$ 24 hour long forecasts $(y_{25:48}^{(i)})_{i=1}^N$. We compute the Root Mean Squared Error (RMSE) between the observations and the average of the forecasts: $\text{RMSE}^2 = T^{-1} \sum_{k=1}^T (Y_k - N^{-1} \sum_{i=1}^N y_k^{(i)})^2$. Additionally, we evaluate the Prediction Interval Coverage Probability (PICP, see [42]) which measures the ratio of observations falling between a 95% confidence interval: $\text{PICP} = T^{-1} \sum_{k=1}^T \mathbb{1}_{[y_L^k, y_U^k]}(Y_k)$, where y_U^k (resp. y_L^k) is the upper (resp. lower) bound of the confidence intervals. Both criteria are reported in Table I.

For our proposed model, predictions can be performed by approximating the predictive density $p_{\theta, \varphi}(y_{k+1} | U_{1:k+1}, Y_{1:k})$ by

$$p_{\hat{\theta}, \hat{\varphi}}^N(y_{k+1}) = \sum_{i=1}^N \omega_k^i p_{\hat{\theta}, \hat{\varphi}}^i(y_{k+1} | \xi_k^i, U_{k+1}),$$

where $p_{\hat{\theta}, \hat{\varphi}}^i(y_{k+1} | \xi_k^i, U_{k+1})$ is the predictive distribution of Y_{k+1} described in Section II-B. In order to explore longer ranges, we run our model to get N samples for any time horizon. The associated intervals containing 95% of the samples are displayed in Figure 2, for 24 hours forecasts.

We compared our model with MC Dropout methods, by implementing recurrent dropout layers as described in [18]. The optimal dropout rate $p_{\text{drop}} = 0.01$ that we tuned by grid search is smaller than the proposed value in the original paper, which may be due to our much longer time series, similarly to results presented in [43]. Additionally, we evaluate the model with $p_{\text{drop}} = 0.05$, which we show slightly degrades performances. The training procedure is similar to traditional

recurrent models ; during inference, we draw 100 samples from the dropout layers, and compute the same average forecasts and intervals as for our model. Despite being based on the same deep learning architecture, the MC dropout model is still largely overconfident, while our proposed model provide more credible empirical confidence intervals. We also experimented with a Gaussian linear Hidden Markov Model (HMM) whose parameters are estimated with the Kalman smoother using the Expectation Maximization (EM, [44]) algorithm. Out of a range of possible latent dimension sizes $d_{\text{hmm}} \in \{1, 2, 4, 6\}$, we selected $d_{\text{hmm}} = 4$ as it yielded the best performances.

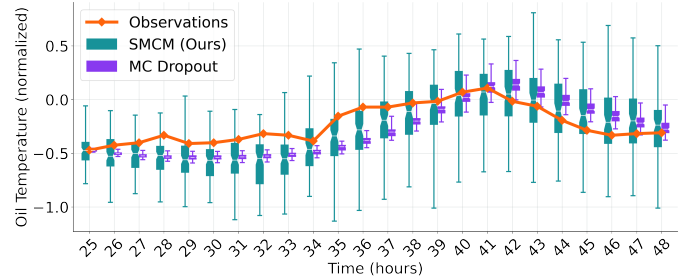


Fig. 2. Forecasting of oil temperature given observations in the lookback window ($k < 24$). Since resampling of particles is no longer available at that point, the uncertainty grows for our model. As a comparison, we plotted the confidence intervals produced by the MC Dropout model (for $p_{\text{drop}} = 0.01$).

TABLE I

COMPARISON OF RMSE, PICP AND COMPUTATION TIME OF OUR MODEL AGAINST THE BENCHMARKED MC DROPOUT METHODS AND HMM. TWO VERSIONS OF THE DROPOUT MODEL WERE EVALUATED, WITH DROPOUT VALUES $p_{\text{drop}} = 0.05$ AND $p_{\text{drop}} = 0.01$. MEAN VALUES OF THE ESTIMATORS ARE DISPLAYED ALONG WITH THEIR VARIANCE.

| | RMSE | PICP | Computation time |
|----------------|-----------------|------|-------------------------|
| SMCL (ours) | 0.24 ± 0.13 | 98% | $210\text{ms} \pm 75.7$ |
| MCD $p = 0.01$ | 0.25 ± 0.15 | 59% | $193\text{ms} \pm 60.4$ |
| MCD $p = 0.05$ | 0.28 ± 0.15 | 65% | $193\text{ms} \pm 60.4$ |
| HMM | 0.44 ± 0.13 | 85% | $994\text{ms} \pm 21.4$ |

V. CONCLUSION

In this paper, we introduced a decoupled architecture for uncertainty estimation on a time series dataset. Our deep neural network backbone is responsible for extracting high level features, while particle filtering in the last layer allows modelling recurrent nonlinear uncertainty. Our proposed model does not suffer from the overconfidence of MC Dropout methods, while significantly improving on the performances of Hidden Markov Models.

We demonstrate the potential behind implementing latent space models as a modified RNN cell ; more complex architectures, such as the GRU network used in the input model, or LSTM cells, could be considered.

Our decoupled architecture enables incorporating uncertainty estimation to an already trained network. This opens the door to multiple, cheap finetuning of last layers parameters, from a global pretraining.

REFERENCES

- [1] Y. Bengio, P. Y. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE transactions on neural networks*, vol. 5 2, pp. 157–66, 1994.
- [2] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, pp. 1735–1780, 1997.
- [3] A. Graves, A. rahman Mohamed, and G. E. Hinton, "Speech recognition with deep recurrent neural networks," 2013, pp. 6645–6649.
- [4] C. S. Crowson, E. J. Atkinson, and T. M. Therneau, "Assessing calibration of prognostic risk scores," *Statistical Methods in Medical Research*, vol. 25, no. 4, pp. 1692–1706, 2016, pMID: 23907781.
- [5] L. H. Mervin, S. Johansson, E. Semanova, K. A. GIBLIN, and O. Engkvist, "Uncertainty quantification in drug design." *Drug discovery today*, vol. 26, no. 2, pp. 474–489, 2021.
- [6] G. E. Hinton and R. Neal, "Bayesian learning for neural networks," 1995.
- [7] M. I. Jordan, Z. Ghahramani, T. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine Learning*, vol. 37, pp. 183–233, 2004.
- [8] M. Fraccaro, S. r. K. Sø nderby, U. Paquet, and O. Winther, "Sequential neural models with stochastic layers," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [9] M. Fortunato, C. Blundell, and O. Vinyals, "Bayesian recurrent neural networks," *arXiv preprint arXiv:1704.02798*, 2017.
- [10] M. Cohen, M. Charbit, and S. Le Corff, "Variational latent discrete representation for time series modelling," *2023 IEEE Statistical Signal Processing Workshop (SSP)*, 2023.
- [11] C. J. Maddison, J. Lawson, G. Tucker, N. Heess, M. Norouzi, A. Mnih, A. Doucet, and Y. Teh, "Filtering variational objectives," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [12] C. Naesseth, S. Linderman, R. Ranganath, and D. Blei, "Variational Sequential Monte Carlo," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Storkey and F. Perez-Cruz, Eds., vol. 84. PMLR, 09–11 Apr 2018, pp. 968–977.
- [13] A. Martin, C. Ollion, F. Strub, S. L. Corff, and O. Pietquin, "The monte carlo transformer: a stochastic self-attention model for sequence prediction," *ArXiv*, vol. abs/2007.08620, 2020.
- [14] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural networks," in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ser. ICML'15. JMLR.org, 2015, p. 1613–1622.
- [15] J. M. Hernandez-Lobato and R. Adams, "Probabilistic backpropagation for scalable learning of bayesian neural networks," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 1861–1869.
- [16] M. Teye, H. Azizpour, and K. Smith, "Bayesian uncertainty estimation for batch normalized deep networks," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 4907–4916.
- [17] A. Foong, D. Burt, Y. Li, and R. Turner, "On the expressiveness of approximate inference in bayesian neural networks," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 15 897–15 908.
- [18] Y. Gal and Z. Ghahramani, "A theoretically grounded application of dropout in recurrent neural networks," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [19] T. Vandal, M. Livingston, C. Pihou, and S. Zimmerman, "Prediction and uncertainty quantification of daily airport flight delays," in *Proceedings of The 4th International Conference on Predictive Applications and APIs*, ser. Proceedings of Machine Learning Research, vol. 82, 24–25 Oct 2018, pp. 45–51.
- [20] M. Wen and E. Tadmor, "Uncertainty quantification in molecular simulations with dropout neural network potentials," *npj Computational Materials*, vol. 6, pp. 1–10, 2020.
- [21] T. Pearce, A. Brintrup, M. Zaki, and A. Neely, "High-quality prediction intervals for deep learning: A distribution-free, ensembled approach," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80, 10–15 Jul 2018, pp. 4075–4084.
- [22] A. Ashukha, A. Lyzhov, D. Molchanov, and D. Vetrov, "Pitfalls of in-domain uncertainty estimation and ensembling in deep learning," *arXiv preprint arXiv:2002.06470*, 2020.
- [23] A. Foong, D. Burt, Y. Li, and R. Turner, "On the expressiveness of approximate inference in bayesian neural networks," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 15 897–15 908.
- [24] N. Brosse, C. Riquelme, A. Martin, S. Gelly, and É. Moulines, "On last-layer algorithms for classification: Decoupling representation from uncertainty estimation," *ArXiv*, vol. abs/2001.08049, 2020.
- [25] R. Douc, É. Moulines, and D. Stoffer, *Nonlinear Time Series: Theory, Methods and Applications with R Examples*. Chapman & Hall, 2013.
- [26] N. J. Gordon, D. J. Salmond, and A. F. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," in *IEE Proceedings F*, vol. 140, no. 2. IET, 1993, pp. 107–113.
- [27] M. K. Pitt and N. Shephard, "Filtering via simulation: Auxiliary particle filters," *J. Amer. Statist. Assoc.*, vol. 94, no. 446, pp. 590–599, 1999.
- [28] R. Douc and O. Cappé, "Comparison of resampling schemes for particle filtering," in *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005*. IEEE, 2005, pp. 64–69.
- [29] N. Chopin and O. Papaspiliopoulos, *An Introduction to Sequential Monte Carlo*. Springer International Publishing, 2020.
- [30] V. Elvira, J. Míguez, and P. M. Djurić, "Adapting the number of particles in sequential monte carlo methods through an online scheme for convergence assessment," *IEEE Transactions on Signal Processing*, vol. 65, no. 7, pp. 1781–1794, 2016.
- [31] V. Elvira, J. Míguez, and P. M. Djurić, "On the performance of particle filters with adaptive number of particles," *Statistics and Computing*, vol. 31, no. 6, pp. 1–18, 2021.
- [32] G. Kitagawa, "Monte carlo filter and smoother for non-gaussian nonlinear state space models," *Journal of Computational and Graphical Statistics*, vol. 5, no. 1, pp. 1–25, 1996.
- [33] C. Andrieu, A. Doucet, and V. Tadic, "On-line parameter estimation in general state-space models," in *Proceedings of the 44th IEEE Conference on Decision and Control*, 2005, pp. 332–337.
- [34] A. Doucet, S. J. Godsill, and C. Andrieu, "On sequential monte carlo sampling methods for bayesian filtering," *Statistics and Computing*, vol. 10, pp. 197–208, 2000.
- [35] S. J. Godsill, A. Doucet, and M. A. West, "Monte carlo smoothing for nonlinear time series," *Journal of the American Statistical Association*, vol. 99, pp. 156 – 168, 2004.
- [36] J. Olsson and J. Westerborn, "Efficient particle-based online smoothing in general hidden markov models : the paris algorithm," *Bernoulli*, vol. 23, pp. 1951–1996, 2014.
- [37] P. Gloaguen, S. L. Corff, and J. Olsson, "A pseudo-marginal sequential Monte Carlo online smoothing algorithm," *Bernoulli*, vol. 28, no. 4, pp. 2606 – 2633, 2022.
- [38] P. Del Moral, A. Doucet, and S. S. Singh, "A backward interpretation of Feynman–Kac formulae," *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 44, pp. 947–975, 2010.
- [39] A. Martin, M.-P. Etienne, P. Gloaguen, S. L. Corff, and J. Olsson, "Backward importance sampling for online estimation of state space models," *arXiv preprint arXiv:2002.05438*, 2020.
- [40] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, pp. 11 106–11 115, May 2021.
- [41] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2015.
- [42] D. L. Shrestha and D. P. Solomatine, "Machine learning approaches for estimation of prediction interval for the model output," *Neural Networks*, vol. 19, no. 2, pp. 225–235, 2006, earth Sciences and Environmental Applications of Computational Intelligence.
- [43] L. Zhu and N. P. Laptev, "Deep and confident prediction for time series at uber," *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 103–110, 2017.
- [44] 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*, vol. 39, pp. 1–38, 1977.