Switching Dynamical Systems with Deep Neural Networks

Ojeda, C.[§], Georgiev, B.[†], Cvejoski, K.[†][?], Schuecker, J.[?], Bauckhage, C.[?], Sánchez, R.^{†‡},

[†]Competence Center Machine Learning Rhine-Ruhr [‡] University of Bonn [?]Fraunhofer Center for Machine Learning and Fraunhofer IAIS [§]Berlin Center for Machine Learning and TU Berlin

Abstract

The problem of uncovering different dynamical regimes is of pivotal importance in time series analysis. Switching dynamical systems provide a solution for modeling physical phenomena whose time series data exhibit different dynamical modes. In this work we propose a novel variational RNN model for switching dynamics allowing for both non-Markovian and non-linear dynamical behavior between and within dynamic modes. Attention mechanisms are provided to inform the switching distribution. We evaluate our model on synthetic and empirical datasets of diverse nature and successfully uncover different dynamical regimes and predict the switching dynamics.

Motivation

Switching linear dynamical system models aim to capture complex (non-linear) time series behaviour via a collection of so-called dynamical modes, each of which is approximated by a linear model [1, 2]. In short, one assumes that at each time step t there is a corresponding categorical latent

Mode Regularization

We enforce dynamical diversity by imposing cost functions to be trained along side the maximization of the lower bound

$$\mathcal{H}[\rho] = -\mathbb{E}_{p_D(\mathsf{x})} \sum_{k=1}^{\mathcal{K}} \tilde{\rho}_k \log \tilde{\rho}_k, \tag{5}$$

where $\tilde{\rho}_k$ is the time-average posterior class probabilities. We maximize then $\mathcal{L}'[q] = \mathcal{L}[q] + \lambda_e \mathcal{H}[\rho]$, where λ_e is a hyperparameter.

Experiments

Lorenz Attractor It is defined by a coupled system of non-linear equations

$$=\sigma(y-x),$$



(6)

(7)

state z_t taking one of K different values and following the Markovian transitions

$$z_{t+1} \mid z_t \sim \pi_{z_t}, \tag{1}$$

where $\pi_{z_t} \in [0, 1]^K$ gives the usual Markov transition probabilities. The classical approach [1, 2] also introduces the continuous latent states $h_t \in \mathbb{R}^p$ — these follow affine dynamics, with the different modes being indexed by z_t ,

$$h_{t+1} = A_{z_{t+1}}h_t + b_{z_{t+1}} + v_t, \quad v_t \sim \mathcal{N}(0, Q_{z_{t+1}}),$$
 (2)

where A_k, Q_k are matrices of the form $\mathbb{R}^{p \times p}$ whereas $b_k \in \mathbb{R}^p$ and $k \in (1, ..., K)$. At last, the observed data points $x_t \in \mathbb{R}^d$ are obtained via

$$\mathbf{x}_t = \mathsf{C}_{z_t} \mathsf{h}_t + \mathsf{d}_{z_t} + \mathsf{w}_t, \quad \mathsf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathsf{S}_{z_{t+1}}), \tag{3}$$

with $C_k \in \mathbb{R}^{d \times p}$, $S_k \in \mathbb{R}^{d \times d}$ and the drift terms $d_k \in \mathbb{R}^d$.

Neural Variational Switching Dynamical Systems

Generative Model

Switching Dynamics (LSTM) h^s_t = f_{θs} (x_t, h^s_{t-1})
 Switching Probabilities π^k_t = softmax [g_{θk} (h^s_{t-1})]
 Switching Prior p(z_t) = Π^K_{k=1} (π^k_t)<sup>z^k_t</sub>
 Emission Parameters [μ^k_t, σ^k_t] = g_{φk} (h^k_{t-1})
 Emission Probabilities p(x^k_t | x^k_{<t}) = N(μ^k_t, diag [(σ^k_t)²])
</sup>



It is defined as:

 $f(t) = H(\cos(\omega_s t) < 0) \cos(\omega_1 t) + [1 - H(\cos(\omega_s t) < 0)] \cos(\omega_2 t)$ Where H is the Heaviside function and $\omega_s < \omega_2 < \omega_1$

Modes Dynamics $h_t^k = f_{\varphi_k} \left(\mathbf{x}_t, \mathbf{h}_{t-1}^k \right)$

Joint Distribution

$$p(\boldsymbol{z}_{\leq T}, \boldsymbol{x}_{\leq T}) = \prod_{k=1}^{K} \prod_{t=1}^{T} \left(\pi_{t}^{k} p\left(\mathbf{x}_{t}^{k} \big| \mathbf{x}_{< t}^{k} \right) \right)^{z_{t}^{k}}.$$



Architecture of the NVSDS model. x_t the data is fed into the recurrent network modeling the modes dynamics and switching dynamics (upper part). The representations h obtained by the experts' dynamics are fed into an MLPs parametrizing a Gaussian distribution for the outputs \hat{x}_t .

We asses these observations quantitatively by defining a binary target vector $\hat{\rho}$ indicating the dynamical mode which is present at a particular time-step. Evaluating the mean squared error between the predicted ρ and $\hat{\rho}$ yields a dissection error, which we average over multiple trials (different initial conditions). The NVSDS (0.09) and the NVSDS-EM (0.1) clearly outperform R-k-Means (0.26), rSLDS (0.4) and MoE (0.47).

Handwrittig

had she slightest effect. Nor is 20 mnutes of discussion is believed 1 activits of Nhrumak's Conven her. " My darling the says son A Property Mar Mar Mar

Dissection of a handwriting signal for the NVSDS model for different sequences. The lower row shows particular letters from the complete sequences for easier comparison.

Conclusion

In the present work we have provided a neural network solution to the problem of switching dynamical systems (SDS).

Inference

Approximate Posterior q(z_t|x_{≤t}) = Π^M_{k=1} (ρ^k_t)^{z^k}
Attention Representation u_k = σ (W_k H_t + V_k h^s_t + b_k)
Prediction and Hidden States H_t = ([x¹_t, h¹_t], [x²_t, h²_t], ..., [x^K_t, h^K_t])
Modes POsterior Probabilities ρ^t_k = softmax [u_k · c_k] ELBO

$$\mathcal{L}[q] = \mathbb{E}_{p_D(\mathsf{x})} \left[\sum_{k=1}^{K} \sum_{t=1}^{T} \left\{ \rho_t^k \log p(\mathsf{x}_t^k | \mathsf{x}_{< t}^k) + \rho_t^k \log \left[\frac{\pi_t^k}{\rho_t^k} \right] \right\} \right], \quad (4)$$

- We build upon variational approximate inference for the categorical variables indexing of the dynamical modes.
- ► We incorporate an attention mechanism for the switching procedure.
- ► We incorporate an entropy regularizer to improve the detection of the modes.

References

- [1] S. Linderman, M. Johnson, A. Miller, R. Adams, D. Blei, and L. Paninski, "Bayesian learning and inference in recurrent switching linear dynamical systems," in *Artificial Intelligence and Statistics*, pp. 914–922, 2017.
- [2] G. Ackerson and K.-S. Fu, "On state estima-tion in switching environments," *IEEE Transactionson Automatic Control*, vol. 15, no. 1, pp. 10–17, 1970.