Regime Learning Particle Filtering

John-Jo Brady, Yunpeng Li

Computer Science Research Centre
University of Surrey, UK

January, 2024

Background: Particle Filters

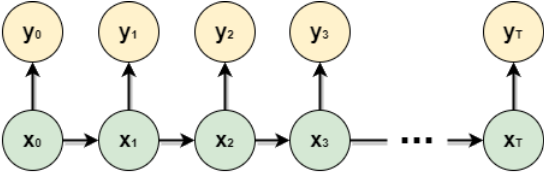

- Set of algorithms to perform inference on state-space models.
- Goal: estimate $\mathbb{E}\left[x_t | y_{0:t}\right]$
- Forms a particle estimate of $P\left(x_t | y_{0:t}\right)$ via sequential Monte Carlo sampling

Background: Particle Filters Cont.

$$P\left(x_t|y_{0:t}\right) = \int P\left(x_{t-1}|y_{0:t-1}\right) \cdot \frac{P\left(x_t|x_{t-1}\right) P\left(y_t|x_t\right)}{P\left(y_t\right)} dx_{t-1}$$

Algorithm 1 Particle Filter Main Loop

1: **for** $t = 1$ to T **do**
2: Sample x_t^n from proposal $M\left(x_t^n|x_{t-1}^n, y_t\right)$
3: Calculate importance weights $w_t^n \propto P\left(y_t|x_t^n\right) \frac{P\left(x_t^n|x_{t-1}^n\right)}{M\left(x_t^n|x_{t-1}^n, y_t\right)}$
4: Auto-normalise $W_t^n = \frac{w_t^n}{\sum_{m=1}^N w_t^m}$
5: Resample x_t^n with probability W_t^n
6: **end for**

Background: Differentiable Particle Filters

- Parameterise (some part of) the SSM using Neural Networks
- Optimised with (autograd) stochastic gradient descent
- In a supervised setting[1], we directly train the MSE between our predicted mean and the ground truth

[1] R. Jonschkowski et al., "Differentiable Particle Filters: End-to-End Learning with Algorithmic Priors", 2018

Background: Regime Switching Particle Filters

- Y. El-Laham et al.[2] introduced a new framework where the system of interest consisted of a discrete set of K SSMs that the system could switch between
- The model probabilities were allowed to depend arbitrarily on the models at all previous time steps
- W. Li et al.[3] made a simple extension where the individual SSMs were learned but the switching dynamic was provided

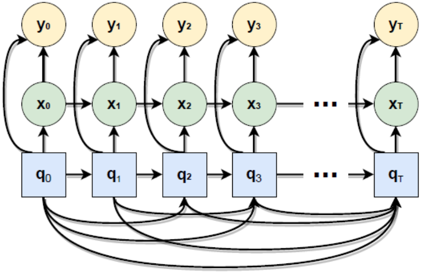

[2] Y. El-Laham et al. "Particle Filtering Under General Regime Switching", 2020

[3] W. Li et al. "Differentiable Bootstrap Particle Filters for Regime-Switching Models", 2023

Learning the Regime Switching Dynamic

- Instead of an arbitrary dependence, we require the previous model indices be encoded in a R dimensional feature vector

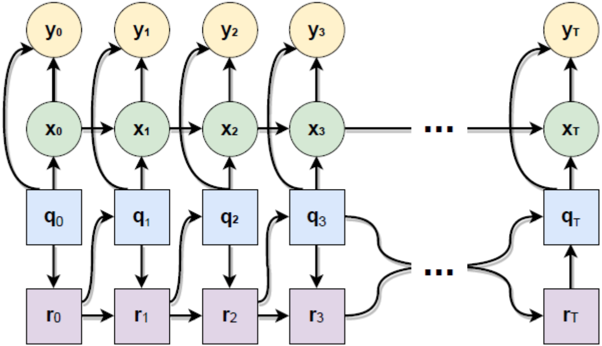

Regime Switching Network Architecture

- It's important that the Markov process (q, r) forgets its past in some sense[4]
- This and the recurrent structure of the model led us to take inspiration from LSTMs

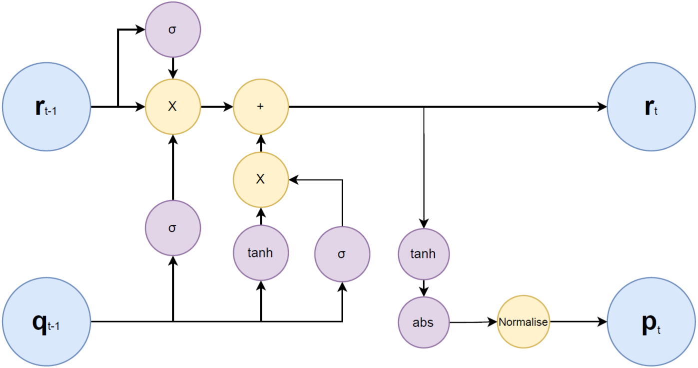

[4] N. Chopin and S. Papaspiliopoulos, "An Introduction to Sequential Monte Carlo", 2020

Algorithm

- To be able to pass gradients through the regime switching network, the probabilities should not be sampled from
- Instead we sample from the models uniformly and (importance) weight the samples by the model probability
- This proposal strategy was proposed by Y. El-Laham et al. to maintain model diversity
- Otherwise use vanilla (bootstrap) particle filtering
- Note: The need for the model index to be discreet necessitates non-differentiable resampling schemes

Experiment Set-up

- We repeat the setting from Y. El-Laham et al. and W. Li et al.
- $x_t | x_{t-1} \sim \mathcal{N}\left(a_q * x_{t-1} + b_q, \sigma_x^2\right)$
- $y_t | x_t \sim \mathcal{N}\left(c_q * \sqrt{|x_t|} + d_q, \sigma_x^2\right)$
- 1st setting uses a Markov Switching dynamic with a constant transition matrix, and the 2nd uses a Polya-Urn distribution

Consequences of Uniform Proposal

- In Y. El-Laham et al. it is stated that a uniform proposal remedies model diversity issues
- More precisely it ensures model diversity at the current time-step

Consequences of Uniform Proposal

- In Y. El-Laham et al. it is stated that a uniform proposal remedies model diversity issues
- More precisely it ensures model diversity at the current time-step

The difference in variance of the weights at the current time-step between uniform and bootstrap proposals:

$$\mathsf{Var}\left[w_t^n\right]_{\mathsf{uniform}} - \mathsf{Var}\left[w_t^n\right]_{\mathsf{bootstrap}}$$

$$= \sum_{i=1}^{K} \left(Q_t\left(i|r_{t-1}\right) K - 1\right) Q_t\left(i|r_{t-1}\right) \mathbb{E}_{x_t}\left[\left(G_t\left(y_t|x_t, i\right)\right)^2\right]$$

Consequences of Uniform Proposal cont.

Assume that the likelihood is informative:

$$\mathsf{Var}\left[w_t^n\right]_{\mathsf{uniform}} - \mathsf{Var}\left[w_t^n\right]_{\mathsf{bootstrap}}$$
$$\approx \left(Q_t\left(j|r_{t-1}\right)K - 1\right)Q_t\left(j|r_{t-1}\right)\mathbb{E}_{x_t}\left[\left(G_t\left(y_t|x_t, j\right)\right)^2\right]$$

Consequences of Uniform Proposal cont.

Assume that the likelihood is informative:

$$\mathsf{Var}\left[w_t^n\right]_{\mathsf{uniform}} - \mathsf{Var}\left[w_t^n\right]_{\mathsf{bootstrap}}$$
$$\approx \left(Q_t\left(j|r_{t-1}\right)K - 1\right)Q_t\left(j|r_{t-1}\right)\mathbb{E}_{x_t}\left[\left(G_t\left(y_t|x_t,j\right)\right)^2\right]$$

- Then the average difference is $\propto \sum_{j=1}^{K}\left(\left(Q_t\left(j|r_{t-1}\right)\right)^3 K - \left(Q_t\left(j|r_{t-1}\right)\right)^2\right) > 0$
- So on average using a uniform proposal should increase the variance at the current time-step
- Empirical simulation showed that whilst the variance of the weights was greatly increased with a uniform proposal, the variance of the resampled particles was similar to bootstrap

Markov Results

	Mean	Best	Worst	SD
DBPF (baseline)	2.562	1.250	9.6019	1.0834
LSTM (baseline)	**0.9756**	0.4935	6.289	0.4766
RLDBPF (ours)	0.9997	**0.4712**	**5.977**	**0.4627**
RSDBPF[5] (oracle)	1.127	0.5277	7.986	0.6174
RSPF[6] (oracle)	0.4701	0.2686	2.5332	0.2041

[5] W. Li et al. "Differentiable Bootstrap Particle Filters for Regime-Switching Models", 2023

[6] Y. El-Laham et al. "Particle Filtering Under General Regime Switching", 2020

Polya-Urn Results

	Mean	Best	Worst	SD
DBPF (baseline)	1.351	0.6444	4.811	0.5264
LSTM (baseline)	1.097	0.5385	**3.807**	**0.3975**
RLDBPF (ours)	**0.8915**	**0.3405**	4.113	0.4130
RSDBPF[7] (oracle)	0.8804	0.4314	3.321	0.3335
RSPF[8] (oracle)	0.6514	0.3105	2.3440	0.2106

[7] W. Li et al. "Differentiable Bootstrap Particle Filters for Regime-Switching Models", 2023

[8] Y. El-Laham et al. "Particle Filtering Under General Regime Switching", 2020

A More General Issue

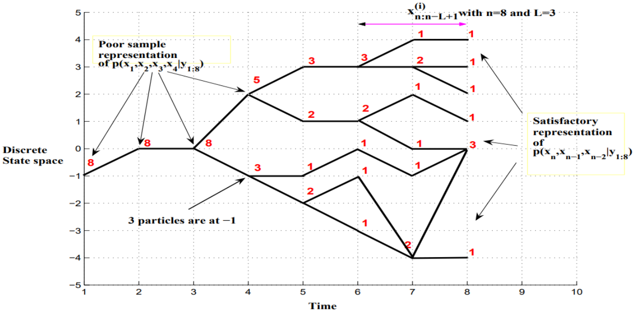

9

- In classical PF literature there is much discussion given to path degeneracy, but little in the current DPF literature

[9] C. Andrieu, A. Doucet, and V. Tadic, "On-line parameter estimation in general state-space models," 2006

A More General Issue

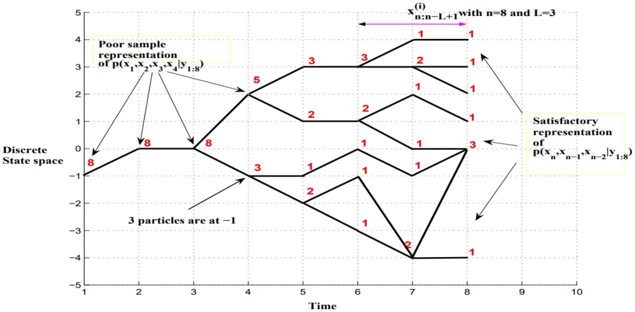

9

- In classical PF literature there is much discussion given to path degeneracy, but little in the current DPF literature
- If we simply auto-grad back through the algorithm then the dependence of late-time results on early-time computation will only be through a small number of particles

[9] C. Andrieu, A. Doucet, and V. Tadic, "On-line parameter estimation in general state-space models," 2006

A More General Issue

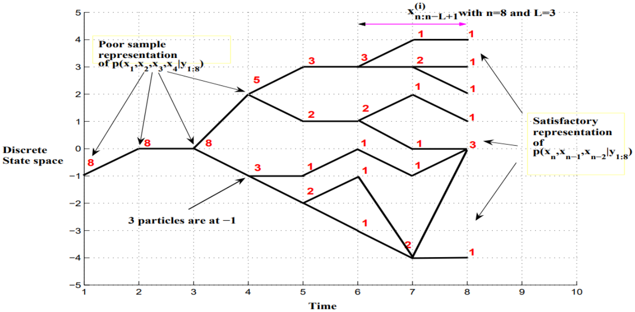

[9]

- In classical PF literature there is much discussion given to path degeneracy, but little in the current DPF literature
- If we simply auto-grad back through the algorithm then the dependence of late-time results on early-time computation will only be through a small number of particles
- This leads to highly variant gradient estimates

[9] C. Andrieu, A. Doucet, and V. Tadic, "On-line parameter estimation in general state-space models," 2006

Where next?

- More complex problems

Where next?

- More complex problems
- Real world data (google smartphone tracking)

Where next?

- More complex problems
- Real world data (google smartphone tracking)
- More constrained scenarios (e.g. allowing the regime to only switch rarely)

Where next?

- More complex problems
- Real world data (google smartphone tracking)
- More constrained scenarios (e.g. allowing the regime to only switch rarely)
- More granular testing

Where next?

- More complex problems
- Real world data (google smartphone tracking)
- More constrained scenarios (e.g. allowing the regime to only switch rarely)
- More granular testing
- Genealogical diversity in DPFs

Where next?

- More complex problems
- Real world data (google smartphone tracking)
- More constrained scenarios (e.g. allowing the regime to only switch rarely)
- More granular testing
- Genealogical diversity in DPFs
- Smarter proposals (non-bootstrap)

References

- R. Jonschkowski, D. Rastogi, and O. Brock, “Differentiable particle filters: End-to-end learning with algorithmic priors,” *CoRR*, vol. abs/1805.11122, 2018.
- C. Andrieu, A. Doucet, and V. Tadic, “On-line parameter estimation in general state-space models,” pp. 332 – 337, 01 2006.
- N. Chopin and O. Papaspiliopoulos, “An introduction to sequential monte carlo,” 01 2020.
- Y. El-Laham, L. Yang, P. Djuric, and M. Bugallo, “Particle filtering under general regime switching,” pp. 2378–2382, 01 2021.
- W. Li, X. Chen, W. Wang, V. Elvira, and Y. Li, “Differentiable bootstrap particle filters for regime-switching models,” *arxiv*, 2023.