Online Inference in Bayesian Non-Parametric Mixture Models under Small Variance Asymptotics

Ajay Kumar Tanwani [1, 2], Sylvain Calinon [1]

[1] Idiap Research Institute, Switzerland.
[2] Ecole Polytechnique Federale de Lausanne, Switzerland.

Non-Parametric Online Learning

- ❑ **Online Robot Learning:** Adapt mixture models (GMM/ MPPCA/ HMM) online with the streaming movement data
- ❑ **Challenges:** Computational overhead of sampling-based and variational techniques limits the widespread use of Bayesian non-parametric mixture models
- ❑ **Solution:** Small variance asymptotic (SVA) analysis of Bayesian non-parametric mixture models for online learning
- ❑ **Applications:** Semi-autonomous teleoperation, robot learning from humans, motion segmentation, subspace tracking and more ...

Given streaming data $\{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_t\}$ with $\boldsymbol{\xi}_t \in \mathbb{R}^D$, update model parameters $\boldsymbol{\theta}_{t+1}$ upon observation of $\boldsymbol{\xi}_{t+1}$.

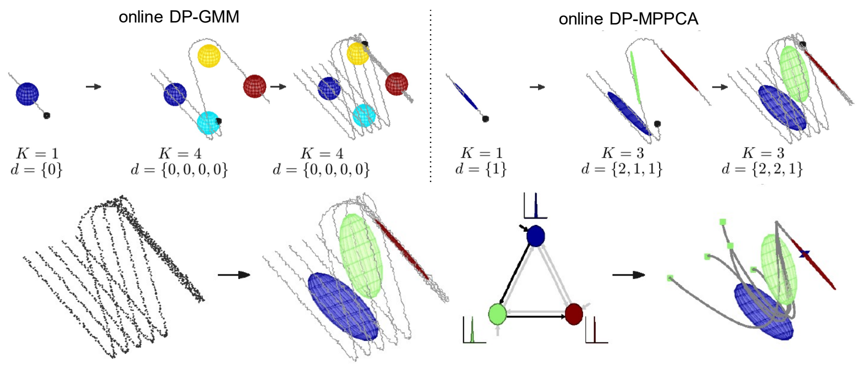

Online Bayesian Non-Parametrics under SVA, $\Sigma_{t,i} \approx \lim_{\sigma^2 \to 0} \sigma^2 \boldsymbol{I}$

Online Dirichlet Process Gaussian Mixture Model (DP-GMM)

Likelihood: $\mathcal{P}(\boldsymbol{\xi}_t|\theta_t) = \sum_{i=1}^{K} \pi_{t,i}\, \mathcal{N}(\boldsymbol{\xi}_{t,i}|\boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i}), \qquad \theta_t = \{\pi_{t,i}, \boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i}\}$

cluster assignment prior: $z_t \sim \mathrm{CRP}(\alpha)$ mean prior: $\boldsymbol{\mu}_{t,i} \sim \mathcal{N}(\boldsymbol{0}, \varrho^2 \boldsymbol{I}_D)$

SVA on the Gibbs sampler yields the online DP-GMM

cluster assignment:

$$z_{t+1} = \arg\min_{j=1:K+1} \begin{cases} \|\boldsymbol{\xi}_{t+1} - \boldsymbol{\mu}_{t,j}\|_2^2, & \text{if } j \leq K \\ \lambda, & \text{otherwise.} \end{cases}$$

parameters update: $z_{t+1} = i$

$$\pi_{t+1,i} = \frac{1}{t+1}\Big(t\pi_{t,i} + 1\Big), \quad \boldsymbol{\mu}_{t+1,i} = \frac{1}{w_{t,i}+1}\Big(w_{t,i}\boldsymbol{\mu}_{t,i} + \boldsymbol{\xi}_{t+1}\Big), \quad w_{t+1,i} = w_{t,i} + 1$$

loss function:

$$\mathcal{L}(z_{t+1}, \boldsymbol{\mu}_{t+1,z_{t+1}}) = \lambda K + \|\boldsymbol{\xi}_{t+1} - \boldsymbol{\mu}_{t+1,z_{t+1}}\|_2^2 \leq \mathcal{L}(z_{t+1}, \boldsymbol{\mu}_{t,z_{t+1}}).$$

Limitation: **Scalable for large scale applications, but cannot encode variance in the data with isotropic Gaussians.**

Solution: **Project the data in low-dimensional subspace and discard the redundant dimensions by SVA in a non-parametric manner.**

Online Dirichlet Process Mixture of Probabilistic Principal Component Analysis (DP-MPPCA)

Likelihood: $\mathcal{P}(\boldsymbol{\xi}_t|\theta_t) = \sum_{i=1}^{K} \pi_{t,i}\, \mathcal{N}(\boldsymbol{\xi}_{t,i}|\boldsymbol{\mu}_{t,i}, \boldsymbol{\Lambda}_{t,i}^{d_{t,i}} \boldsymbol{\Lambda}_{t,i}^{d_{t,i}\top} + \sigma^2 \boldsymbol{I}), \ \theta_t = \{\pi_{t,i}, \boldsymbol{\mu}_{t,i}, d_{t,i}, \boldsymbol{\Lambda}_{t,i}^{d_{t,i}}\}$

hierarchical exponential prior on $\boldsymbol{\Lambda}_{t,i}^{d_{t,i}} \in \mathbb{R}^{D \times d_{t,i}}$, $z_t \sim \mathrm{CRP}(\alpha), \ \boldsymbol{\mu}_{t,i} \sim \mathcal{N}(\boldsymbol{0}, \varrho^2 \boldsymbol{I}_D)$

SVA on the partially collapsed Gibbs sampler yields the online DP-MPPCA

cluster assignment:

$$z_{t+1} = \arg\min_{j=1:K+1} \begin{cases} \left\| (\boldsymbol{\xi}_{t+1} - \boldsymbol{\mu}_{t,j}) - \rho_j \boldsymbol{U}_{t,j}^{d_{t,j}} \boldsymbol{U}_{t,j}^{d_{t,j}\top} (\boldsymbol{\xi}_{t+1} - \boldsymbol{\mu}_{t,j}) \right\|_2^2, & \text{if } j \leq K \\ \lambda, & \text{otherwise,} \end{cases}$$

parameters update: $z_{t+1} = i$

$$\boldsymbol{U}_{t+1,i}^{d_{t,i}} = [\boldsymbol{U}_{t,i}^{d_{t,i}}, \ \tilde{\boldsymbol{p}}_{t+1,i}]\ \boldsymbol{R}_{t+1,i}$$

[solved using eigendecomposition of size $(d_{t,i}+1) \times (d_{t,i}+1)$]

$$d_{t+1,i} = \arg\min_{d=0:D-1} \left\{ \lambda_1 d + \text{weighted average of} \begin{bmatrix} \mathrm{dist}(\boldsymbol{\xi}_{t+1}, \boldsymbol{\mu}_{t+1,i}, \boldsymbol{U}_{t+1,i}^{0})^2 \\ \vdots \\ \mathrm{dist}(\boldsymbol{\xi}_{t+1}, \boldsymbol{\mu}_{t+1,i}, \boldsymbol{U}_{t+1,i}^{d_{t,i}+1})^2 \end{bmatrix} \right\}$$

$$\boldsymbol{\Lambda}_{t+1,i}^{d_{t+1,i}} = \boldsymbol{U}_{t+1,i}^{d_{t+1,i}} \sqrt{\boldsymbol{\Sigma}_{t+1,i}^{(\mathrm{diag})}}, \qquad \boldsymbol{\Sigma}_{t+1,i} = \boldsymbol{\Lambda}_{t+1,i}^{d_{t+1,i}} \boldsymbol{\Lambda}_{t+1,i}^{d_{t+1,i}\top} + \sigma^2 \boldsymbol{I}.$$

loss function:

$$\mathcal{L}(z_{t+1}, d_{t+1,z_{t+1}}, \boldsymbol{\mu}_{t+1,z_{t+1}}, \boldsymbol{U}_{t+1,z_{t+1}}^{d_{t+1,z_{t+1}}}) = \lambda K + \lambda_1 d_{t+1,z_{t+1}} \dots + \mathrm{dist}(\boldsymbol{\xi}_{t+1}, \boldsymbol{\mu}_{t+1,z_{t+1}}, \boldsymbol{U}_{t+1,z_{t+1}}^{d_{t+1,z_{t+1}}})^2 \leq \mathcal{L}(z_{t+1}, d_{t,z_{t+1}}, \boldsymbol{\mu}_{t,z_{t+1}}, \boldsymbol{U}_{t,z_{t+1}}^{d_{t,z_{t+1}}}).$$

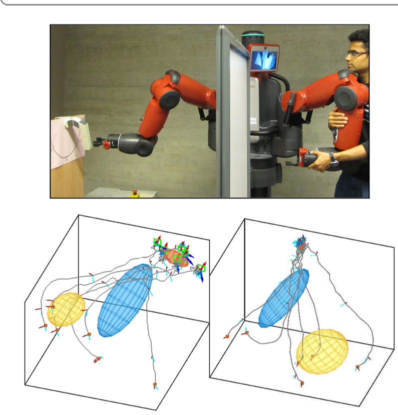

online task-parameterized robot learning

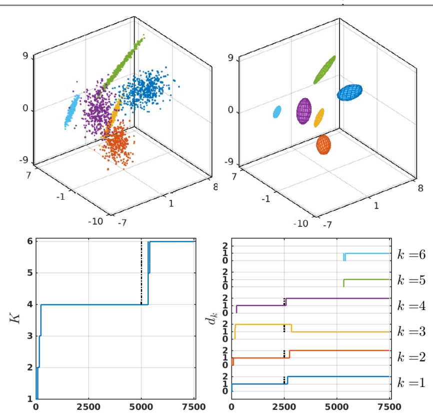

clustering non-stationary data

Summary

- ❑ Scalable non-parametric online learning to adapt the model on the fly with simple deterministic updates
- ❑ Number of clusters and subspace dimension of each cluster adapt with streaming data and the penalty parameters act as regularization terms
- ❑ Temporal patterns are incorporated using online hierarchical Dirichlet process hidden semi-Markov model (HDP-HSMM) [1]
- ❑ Learning the model online from a few human demonstrations is a pragmatic approach to teach new skills to robots

[1] A. K. Tanwani, S. Calinon, "Small Variance Asymptotics for Non-Parametric Online Robot Learning", arXiv:1610.02468, 2016