Recursive Learning of Asymptotic Variational Objectives

Alessandro Mastrototaro Mathias Müller Jimmy Olsson

Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

{alemas, matmul, jimmyol}@kth.se

KTH VETENSKAP OCH KONST

Abstract

We propose the online sequential IWAE (OSIWAE), which maximizes an IWAE-type variational lower bound on the asymptotic contrast function using stochastic approximation. Unlike the original IWAE framework by Burda et al. [1], our approach enables robust online learning for streaming data. [2].

State-Space Models (SSMs)

State-space models (SSMs) are widely used for sequential time-series data, where we model latent (hidden) states $z_1, z_2, \dots, z_T$ and observations $x_1, x_2, \dots, x_T$. These are defined by:

- A transition density $p_\theta(z_t \mid z_{t-1})$ governing the evolution of the latent states, and an emission density $p_\theta(x_t \mid z_t)$ describing the generation of observations.
- The primary objectives: estimating the posterior

$$p_\theta(z_{1:t} \mid x_{1:t}) = \frac{p_\theta(x_{1:t}, z_{1:t})}{p_\theta(x_{1:t})} = \frac{\prod_{t=1}^{T} p_\theta(x_t \mid z_t) p_\theta(z_t \mid z_{t-1})}{p_\theta(x_{1:t})},$$

where $p_\theta(x_{1:t})$ is the marginal likelihood, and learning model parameters θ.

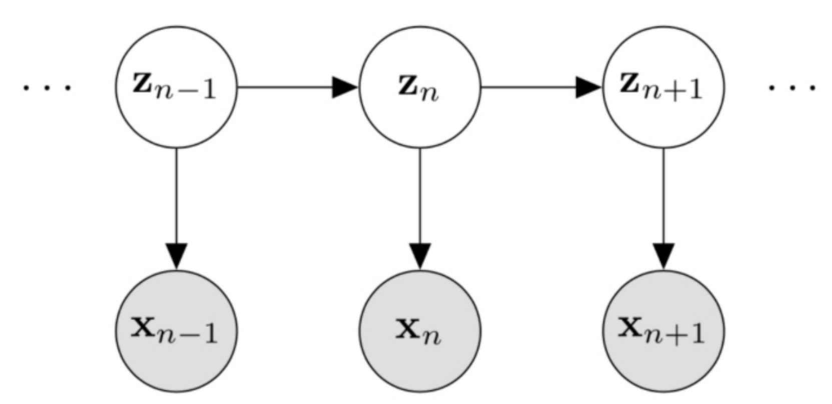

Importance-Weighted Autoencoder (IWAE) ELBO

In Variational Inference (VI) we approximate the intractable posterior $p_\theta(z_{1:t}|x_{1:t})$ using a variational distribution $q_\phi(z_{1:t}|x_{1:t})$. The Importance-Weighted Autoencoder (IWAE) [1] introduces the IWAE-Evidence Lower Bound (ELBO), defined as:

$$\log p_\theta(x_{1:t}) \geq \mathcal{L}_K(\theta, \phi; x) = \mathbb{E}_{z_{1:K} \sim q_\phi}\left[\log\left(\frac{1}{K}\sum_{k=1}^{K}\frac{p_\theta(z_{1:t}^k, x_{1:t})}{q_\phi(z_{1:t}^k | x_{1:t})}\right)\right].$$

- Maximizing $\mathcal{L}_K$ with respect to ϕ brings the approximation closer to the true log-likelihood $\log p_\theta(x_{1:t})$.
- The bound $\mathcal{L}_K$ becomes tighter as the number of samples K increases, improving inference quality.

Online Sequential IWAE (OSIWAE)

We extend the IWAE framework to sequential data by leveraging the asymptotic contrast function, defined as:

$$\ell(\theta) = \lim_{t\to\infty} \frac{1}{t} \log p_\theta(x_{0:t}),$$

which is factorized as:

$$\log p_\theta(x_{0:t}) = \log p_\theta(x_0) + \sum_{s=1}^{t} \log p_\theta(x_s \mid x_{0:s-1}).$$

For each term, $\log p_\theta(x_s|x_{0:s-1}) = V_\theta$, we introduce the Contrast Lower Bound Objective (COLBO), which is a variational lower bound:

$$V_\theta^N(x_{t+1}, \phi_t) = \mathbb{E}_{q_\lambda^{\otimes N}}\left[\log\left(\frac{1}{N}\sum_{i=1}^{N}\frac{m_\theta(Z_t^i|Z_{t-1}^i) g_\theta(x_t|Z_t^i)}{r_\phi(Z_t^i|Z_{t-1}^i, x_t)}\right)\right] \leq V_\theta$$

At each time step, we maximize the COLBO objective to jointly update the model parameters θ and the proposal parameters ϕ.

Gradient Calculation and Parameter Updates

At each time step t, we compute the gradients of the COLBO objective and update the parameters using the following formulations:

$$G_\theta^M(x_{t+1}, \phi_t^\theta, \psi_t^\theta) = \mathbb{E}_{\phi_t^\theta \otimes v}\left[\frac{\sum_{i=1}^{M} \omega_\theta(Z^i, x_{t+1}, U^i) \nabla_\theta \omega_\theta(Z^i, x_{t+1}, U^i)}{\sum_{i'=1}^{M} \omega_\theta(Z^{i'}, x_{t+1}, U^{i'})}\right]$$
$$+ M \log\left(\frac{1}{M}\sum_{i=1}^{M} \omega_\theta(Z^i, x_{t+1}, U^i)\right)\left(\phi_t^\theta(X^M) - \mathbb{E}_{\phi_t^\theta}[\phi_t^\theta(X)]\right).$$

Parameter Update: The model parameters θ_t are updated using a Robbins-Monro algorithm:

$$\theta_{t+1} \leftarrow \theta_t + \gamma_t G_\theta^M(y_{t+1}, \phi_t^\theta, \psi_t^\theta),$$

Experiments: Growth and Gaussian Models

Growth Model: This benchmark features nonlinear dynamics and bimodal proposal challenges. The transition density is $m_\theta(z_t \mid z_{t-1}) = N(z_t; a(z_{t-1}), \sigma_u^2)$, where $a(z_{t-1}) = \alpha_0 z_{t-1} + \frac{\alpha_1 z_{t-1}}{1+z_{t-1}^2} + \alpha_2 \cos(1.2t)$, and the emission density is $g_\theta(x_t \mid z_t) = N(x_t; b z_t^2, \sigma_v^2)$.

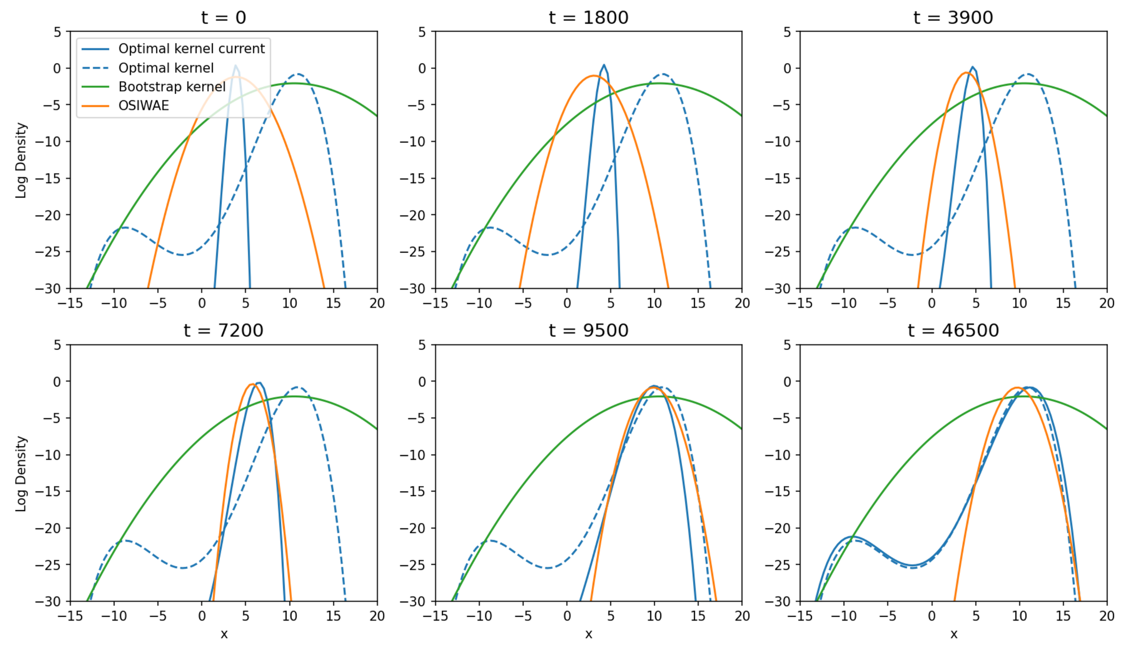

Figure 1. Log-densities of the learned proposal for the Growth Model.

Gaussian Model: This 10-dimensional linear Gaussian system tests the algorithm's efficiency in estimating parameters and latent states. The transition density is $m_\theta(z_t \mid z_{t-1}) = N(z_t; A z_{t-1}, S_u S_u^\top)$, and the emission density is $g_\theta(x_t \mid z_t) = N(x_t; B z_t, S_v S_v^\top)$.

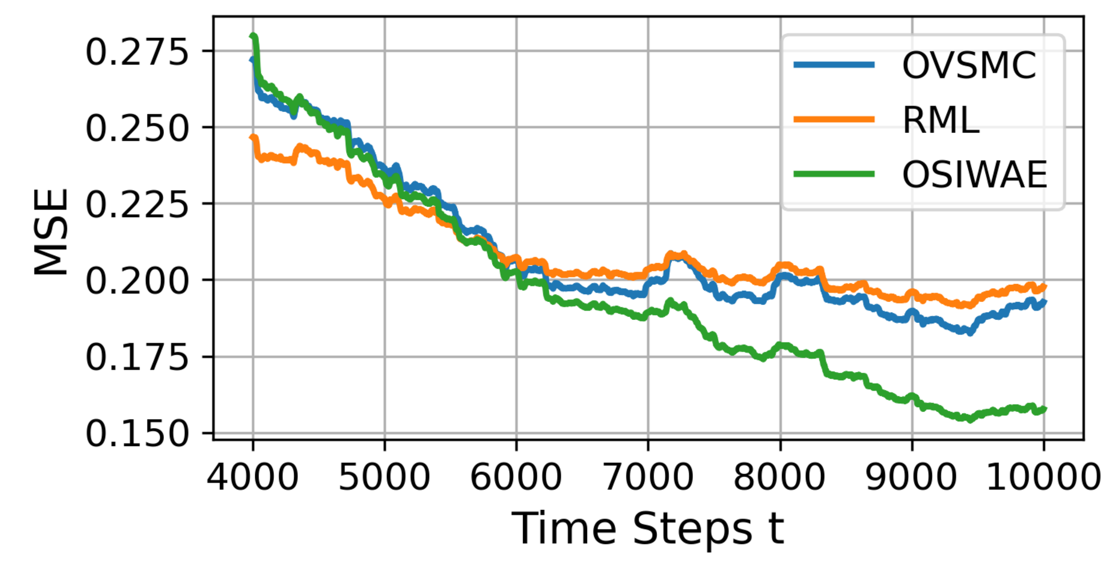

Figure 2. State estimation errors for the Gaussian Model.

Acknowledgments

This work is supported by the Swedish Research Council, grant 2018-05230, and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) *Online learning in dynamical generative models*.

References

[1] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders, 2015.

[2] Alessandro Mastrototaro, Mathias Müller, and Jimmy Olsson. Recursive learning of asymptotic variational objectives, 2024.