

Supplementary Material for: “Efficient and Flexible Inference for Stochastic Systems”

5.1 Update equations for Parameter and State Estimation

Due to the model assumption in [Gorbach et al. \[2017\]](#), the true conditional distributions $p(\boldsymbol{\theta} \mid \mathbf{X}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}, \boldsymbol{\sigma})$ and $p(\mathbf{x}_u \mid \boldsymbol{\theta}, \mathbf{X}_{-u}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}, \boldsymbol{\sigma})$ are Gaussian distributed, where \mathbf{X}_{-u} denotes all states excluding state \mathbf{x}_u (i.e. $\mathbf{X}_{-u} := \{\mathbf{x} \in \mathbf{X} \mid \mathbf{x} \neq \mathbf{x}_u\}$). For didactical reasons, we write the true conditional distributions in canonical form:

$$\begin{aligned} p(\boldsymbol{\theta} \mid \mathbf{X}, \mathbf{Y}, \boldsymbol{\phi}) &= h(\boldsymbol{\theta}) \times \exp(\boldsymbol{\eta}_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}, \boldsymbol{\sigma})^T \mathbf{t}(\boldsymbol{\theta}) - a_{\boldsymbol{\theta}}(\boldsymbol{\eta}_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}, \boldsymbol{\sigma}))) \\ p(\mathbf{x}_u \mid \boldsymbol{\theta}, \mathbf{X}_{-u}, \mathbf{Y}, \boldsymbol{\phi}) &= h(\mathbf{x}_u) \times \exp(\boldsymbol{\eta}_u(\boldsymbol{\theta}, \mathbf{X}_{-u}, \mathbf{Y}, \boldsymbol{\phi})^T \mathbf{t}(\mathbf{x}_u) - a_u(\boldsymbol{\eta}_u(\mathbf{X}_{-u}, \mathbf{Y}, \boldsymbol{\phi}))) \end{aligned}$$

The particular form of $q(\boldsymbol{\theta} \mid \boldsymbol{\lambda})$ and $q(\mathbf{x}_u \mid \boldsymbol{\psi}_u)$ of the mean-field approximations [14](#), are designed to be in the same exponential family as the true conditional distributions:

$$\begin{aligned} q(\boldsymbol{\theta} \mid \boldsymbol{\lambda}) &:= h(\boldsymbol{\theta}) \exp(\boldsymbol{\lambda}^T \mathbf{t}(\boldsymbol{\theta}) - a_{\boldsymbol{\theta}}(\boldsymbol{\lambda})) \\ q(\mathbf{x}_u \mid \boldsymbol{\psi}_u) &:= h(\mathbf{x}_u) \exp(\boldsymbol{\psi}_u^T \mathbf{t}(\boldsymbol{\psi}_u) - a_u(\boldsymbol{\psi}_u)), \end{aligned}$$

whose optimal variational parameters are given by:

$$\hat{\boldsymbol{\lambda}} := \begin{pmatrix} \mathbb{E}_Q \boldsymbol{\Omega}_{\boldsymbol{\theta}}^{-1} \mathbf{r}_{\boldsymbol{\theta}} \\ -\frac{1}{2} \mathbb{E}_Q \boldsymbol{\Omega}_{\boldsymbol{\theta}}^{-1} \end{pmatrix}, \quad (16)$$

$$\hat{\boldsymbol{\psi}}_u := \begin{pmatrix} \mathbb{E}_Q \boldsymbol{\Omega}_u^{-1} \mathbf{r}_u \\ -\frac{1}{2} \mathbb{E}_Q \boldsymbol{\Omega}_u^{-1} \end{pmatrix} \quad (17)$$

Next we analytically derive the parameters $\mathbf{r}_{\boldsymbol{\theta}}$ and $\boldsymbol{\Omega}_{\boldsymbol{\theta}}$ in equation [16](#) which we obtain from the true conditional distribution over ODE parameters:

$$\begin{aligned} p(\boldsymbol{\theta} \mid \mathbf{X}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}) &\stackrel{(a)}{=} p(\boldsymbol{\theta} \mid \mathbf{X}, \boldsymbol{\phi}, \boldsymbol{\gamma}) \\ &\stackrel{(b)}{=} Z_{\boldsymbol{\theta}}^{-1}(\mathbf{X}) \int p(\dot{\mathbf{X}} \mid \mathbf{X}, \boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\gamma}) p(\dot{\mathbf{X}} \mid \mathbf{X}, \boldsymbol{\phi}) d\dot{\mathbf{X}} \\ &\stackrel{(c)}{=} Z_{\boldsymbol{\theta}}^{-1}(\mathbf{X}) \prod_k \mathcal{N}(\mathbf{f}_k(\mathbf{X}, \boldsymbol{\theta}) \mid \mathbf{m}_k, \boldsymbol{\Lambda}_k^{-1}) \\ &\stackrel{(d)}{=} Z_{\boldsymbol{\theta}}^{-1}(\mathbf{X}) \prod_k \mathcal{N}(\mathbf{B}_{\boldsymbol{\theta}k} \boldsymbol{\theta} + \mathbf{b}_{\boldsymbol{\theta}k} \mid \mathbf{m}_k, \boldsymbol{\Lambda}_k^{-1}) \\ &\stackrel{(e)}{=} Z'_{\boldsymbol{\theta}}{}^{-1}(\mathbf{X}) \prod_k \mathcal{N}(\boldsymbol{\theta} \mid (\mathbf{B}_{\boldsymbol{\theta}k}^T \boldsymbol{\Lambda}_k \mathbf{B}_{\boldsymbol{\theta}k})^{-1} \mathbf{B}_{\boldsymbol{\theta}k}^T \boldsymbol{\Lambda}_k (\mathbf{m}_k - \mathbf{f}_{\boldsymbol{\theta}k}), (\mathbf{B}_{\boldsymbol{\theta}k}^T \boldsymbol{\Lambda}_k \mathbf{B}_{\boldsymbol{\theta}k})^{-1}) \\ &\stackrel{(f)}{=} \mathcal{N}(\boldsymbol{\theta} \mid \mathbf{r}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\boldsymbol{\theta}}), \end{aligned}$$

where $Z_{\boldsymbol{\theta}}(\mathbf{X})$ and $Z'_{\boldsymbol{\theta}}(\mathbf{X})$ normalize the distributions and \mathbf{m}_k and $\boldsymbol{\Lambda}_k$ are defined as above. In (a) we notice that $\boldsymbol{\theta}$ does not directly depend on the observations \mathbf{Y} but instead indirectly through the states \mathbf{X} . In (b) we substitute the product of experts and in (c) we analytically integrate out the state derivatives. We rewrite the ODE \mathbf{f}_k as a linear combination of the ODE parameters (i.e. $\mathbf{B}_{\boldsymbol{\theta}k} \boldsymbol{\theta} + \mathbf{b}_{\boldsymbol{\theta}k} \stackrel{!}{=} \mathbf{f}_k(\mathbf{X}, \boldsymbol{\theta})$) in (d) and in (e) we normalize each factor w.r.t. $\boldsymbol{\theta}$. In (g) we normalize the product of Gaussians where mean and covariance are given by:

$$\mathbf{r}_{\boldsymbol{\theta}} := \boldsymbol{\Omega}_{\boldsymbol{\theta}} \sum_k \mathbf{B}_{\boldsymbol{\theta}k}^T \boldsymbol{\Lambda}_k (\mathbf{m}_k - \mathbf{b}_{\boldsymbol{\theta}k}), \quad \boldsymbol{\Omega}_{\boldsymbol{\theta}}^{-1} := \sum_k \mathbf{B}_{\boldsymbol{\theta}k}^T \boldsymbol{\Lambda}_k \mathbf{B}_{\boldsymbol{\theta}k}.$$

Similarly, we analytically derive the parameters \mathbf{r}_u and $\mathbf{\Omega}_u$ in equation 17 which we obtain from the true conditional distribution over an individual state:

$$\begin{aligned}
p(\mathbf{x}_u \mid \boldsymbol{\theta}, \mathbf{X}_{/\{\mathbf{x}_u\}}, \mathbf{Y}, \boldsymbol{\phi}, \boldsymbol{\gamma}) &= Z_u^{-1}(\boldsymbol{\theta}) \prod_k \mathcal{N}(\mathbf{f}_k(\mathbf{X}, \boldsymbol{\theta}) \mid \mathbf{m}_k, \boldsymbol{\Lambda}_k^{-1}) \mathcal{N}(\mathbf{x}_u \mid \boldsymbol{\mu}_u(\mathbf{Y}), \boldsymbol{\Sigma}_u) \\
&\stackrel{(g)}{=} Z_u^{-1}(\boldsymbol{\theta}) \prod_k \mathcal{N}(\mathbf{B}_{uk}\mathbf{x}_u + \mathbf{b}_{uk} \mid \mathbf{m}_k, \boldsymbol{\Lambda}_k^{-1}) \mathcal{N}(\mathbf{x}_u \mid \boldsymbol{\mu}_u(\mathbf{Y}), \boldsymbol{\Sigma}_u) \\
&= Z_u'^{-1}(\boldsymbol{\theta}) \prod_k \mathcal{N}\left(\mathbf{x}_u \mid (\mathbf{B}_{uk}^T \boldsymbol{\Lambda}_k \mathbf{B}_{uk})^{-1} \mathbf{B}_{uk}^T \boldsymbol{\Lambda}_k (\mathbf{m}_k - \mathbf{b}_{uk}), (\mathbf{B}_{uk}^T \boldsymbol{\Lambda}_k \mathbf{B}_{uk})^{-1}\right) \\
&\quad \mathcal{N}(\mathbf{x}_u \mid \boldsymbol{\mu}_u(\mathbf{Y}), \boldsymbol{\Sigma}_u) \\
&= \mathcal{N}(\mathbf{x}_u \mid \mathbf{r}_u, \mathbf{\Omega}_u),
\end{aligned}$$

where $Z_u(\boldsymbol{\theta})$ and $Z_u'(\boldsymbol{\theta})$ normalize the distributions. For a partially observed system, the mean $\boldsymbol{\mu}_u(\mathbf{Y})$ and covariance $\boldsymbol{\Sigma}_u$ are given by $\boldsymbol{\mu}_u(\mathbf{Y}) := \sigma^{-2} (\sigma^{-2} \mathbf{A}^T \mathbf{A} + \mathbf{C}_\phi)^{-1} \mathbf{A}^T \mathbf{Y}$ and $\boldsymbol{\Sigma}_u^{-1} := \sigma^{-2} \mathbf{A}^T \mathbf{A} + \mathbf{C}_\phi^{-1}$, with matrix \mathbf{A} accommodating for unobserved states by encoding the linear relationship between latent states and observations (i.e. $\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{E}$, $\mathbf{E} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$). Once more, in (e) we define \mathbf{B}_{uk} and \mathbf{b}_{uk} such that the ODE \mathbf{f}_k is expressed as a linear combination of the state \mathbf{x}_k (i.e. $\mathbf{B}_{uk}\mathbf{x}_u + \mathbf{b}_{uk} \stackrel{!}{=} \mathbf{f}_k(\mathbf{X}, \boldsymbol{\theta})$). The mean and covariance are given by:

$$\mathbf{r}_u := \boldsymbol{\Omega} \left(\sum_k \mathbf{B}_{uk}^T \boldsymbol{\Lambda}_k (\mathbf{m}_k - \mathbf{b}_{uk}) + \boldsymbol{\Sigma}_u^{-1} \boldsymbol{\mu}_u(\mathbf{Y}) \right), \quad \boldsymbol{\Omega}_u^{-1} := \sum_k \mathbf{B}_{uk}^T \boldsymbol{\Lambda}_k \mathbf{B}_{uk} + \boldsymbol{\Sigma}_u^{-1}.$$

5.2 Lorenz Attractor

Parameter and state estimation has previously been shown in the experiments section 4 where we show the phase plots in figure 4. In figure 8 we show the same estimation in the time domain using mean-field gradient matching by Gorbach et al. [2017] for one sample path.

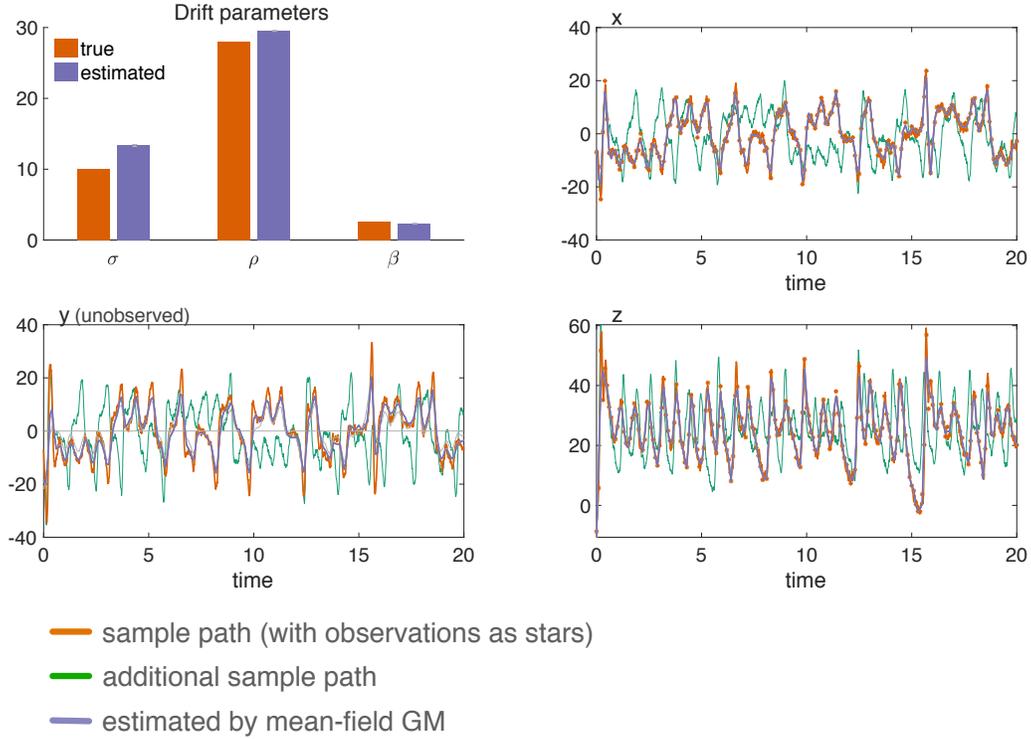


Figure 8: Mean-field gradient matching [Gorbach et al., 2017] is used for parameter and state estimation for the partially observed Lorenz attractor SDE with the y dimension (bottom left) remaining unobserved. Red trajectories respectively denote the sample paths (some of which are observed) and purple denotes their estimation. Both the drift parameters as well as the states and most notably the unobserved state y are estimated well. The experimental setup is the same as the one used by Vrettas et al. [2011]. We included another sample path (green) to illustrate the stochasticity of the system.

5.3 Lorenz 96

Parameter and state estimation for the Lorenz 96 SDE has also been shown previously in the experiments section 4. In figure 9 we additionally show the parameter and state trajectory estimation using mean-field gradient matching by Gorbach et al. [2017] for one sample path.

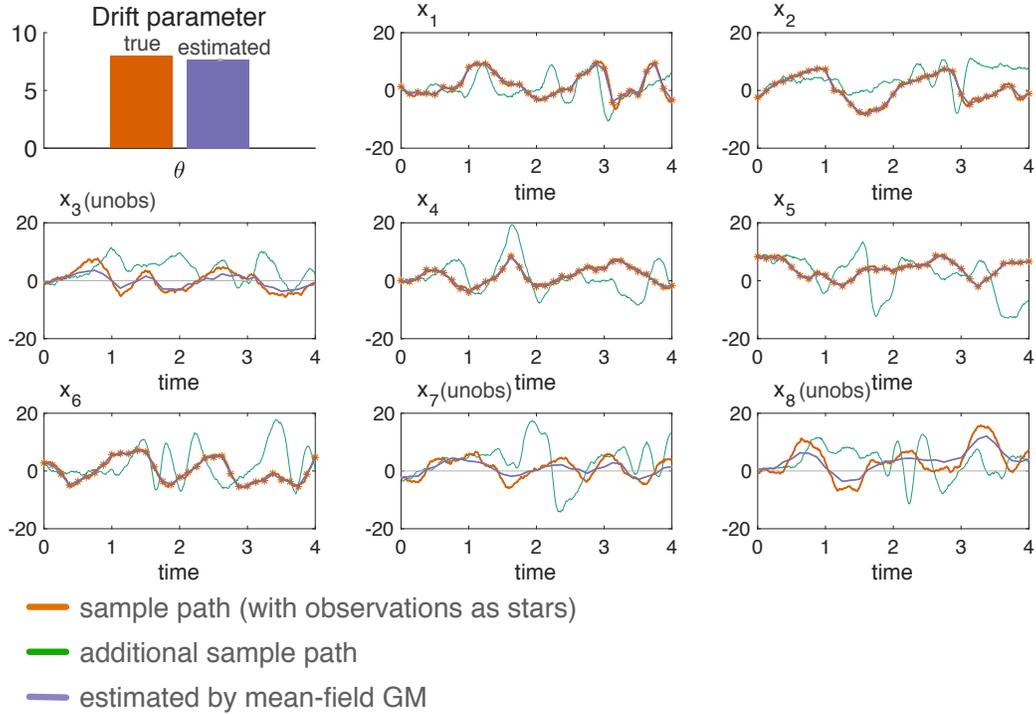


Figure 9: Mean-field gradient matching is used for parameter and state estimation for the partially observed Lorenz 96 SDE. The first 8 of 40 dimensions are shown. Red trajectories respectively denote the sample paths (some of which are observed) and purple denotes their estimation. Both the drift parameters as well as the unobserved state trajectories are estimated well. The experimental setup is the same as the one used by Vrettas et al. [2011]. We included another sample path (green) to illustrate the stochasticity of the system.