Estimating Hidden Dynamics in Stochastic Reaction Networks from Partial Observations

Mingkai Yu [1] Muruhan Rathinam [2]

[1]Melchizedek Education and Research Institute (MEERI) [2]Department of Mathematics and Statistics, University of Maryland, Baltimore County

Abstract

Stochastic reaction network models arise in intracellular reaction systems, as molecules interact through random events. These systems can be modeled as continuous-time Markov chains whose state space represents the count of the species involved. We consider the filtering problem with partial observations: estimating the conditional probability distribution of a hidden state given exact partial state observations in two settings (continous-time and discrete-time observation). We propose weighted Monte Carlo particle filtering algorithms. We provide a rigorous analysis as well as numerical examples to illustrate our method and compare it with alternatives.

Background: Stochastic Reaction Networks

Stochastic Reaction Network: n molecular **species** and m **reaction channels**.

$Z(t) \in \mathbb{Z}_+^n$: the copy number vector of the species at time t. $Z(t)$ is a Continuous-time Markov Chain (CTMC).

- **Stoichiometric vector**: $\nu_j \in \mathbb{Z}^n$, $j = 1, ...m$.
- **Propensity functions** $a_j(z)$:

$$a_j(z)h + o(h) = \text{Probability that reaction } j \text{ will occur in infinitesimal time interval } (t, t+h], \text{ given} Z(t) = z.$$

The time evolution of the probability mass function $p(t,z) = P\{Z(t) = z\}$ is given by the **Kolmogorov's forward equations** also known as the **chemical master equations** (CME):

$$p'(t,z) = \sum_{j=1}^{m} a_j(z-\nu_j)\, p(t, z-\nu_j) - \sum_{j=1}^{m} a_j(z)\, p(t,z), \quad \forall t > 0, z \in \mathbb{Z}_+^n. \tag{1}$$

Example 1: Gene Expression

$$S \xrightarrow{c_1} S + A, \quad \varnothing \xrightarrow{c_2} S, \quad S \xrightarrow{c_3} \varnothing,$$

Here, $Z(t) = (\#A(t), \#S(t))$.

$$\nu = [\nu_1, \nu_2, \nu_3] = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \end{bmatrix}$$

with propensities $a_1(z) = c_1 z_2$, $a_2(z) = c_2$, $a_3(z) = c_3 z_2$.

Example 2: Genetic Toggle Switch

$$\varnothing \xrightarrow{a_1} S_1,\ S_1 \xrightarrow{a_2} \varnothing,\ \varnothing \xrightarrow{a_3} S_2,\ S_2 \xrightarrow{a_4} \varnothing.$$

with $a_1(z) = \frac{\alpha_1}{1+Z_2^\beta}$, $a_2(z) = z_1$, $a_3(z) = \frac{\alpha_2}{1+z_1^\gamma}$, $a_4(z) = z_2$.

Example 3: Pure Death Process

$$S \xrightarrow{c} \varnothing$$

Filtering Problem: Uncovering Hidden Dynamics

Stochastic filtering theory aims to find the conditional distribution of a signal $X(t)$ based on the observation of a related observation $Y(t)$.

- **Partial Observation:** Exact (noiseless) observations of the copy number of the last n_2 species
- $Z(t) = (X(t), Y(t))$, where $X(t) \in \mathbb{Z}_+^{n_1}$ is the **unobserved state** and $Y(t) \in \mathbb{Z}_+^{n_2}$ is the **observed state**.
- **Scenario 1:** Partial Observations in continuous time. Goal is to estimate

$$\pi(t,x) = P\{X(t) = x \,|\, Y(s) = y(s),\ 0 \le s \le t\} \quad \forall x \in \mathbb{Z}_+^{n_1}. \tag{2}$$

- **Scenario 2:** Partial Observations in discrete time snapshots. Goal is to estimate

$$\pi(t,z) = P\{Z(t) = z \,|\, Y(t_i) = y_i\} \quad \forall z \in \mathbb{Z}_+^n,\ t \in [0,T]. \tag{3}$$

Method: Analysis and Simulation

Continuous-time observation: Filtering Equations and Monte Carlo Simulation

$\pi(t,x)$ satisfies filtering equations, a system of nonlinear ODEs with jumps (often computationally infeasible).

We utilize a **weighted Monte Carlo simulation**, also known as a **particle filter**.

$$\hat{\pi}(t,x) = \frac{\sum_{i=1}^{N_s} 1_{\{x\}}(V^{(i)}(t))\, w^{(i)}(t)}{\sum_{i=1}^{N_s} w^{(i)}(t)}. \tag{4}$$

where $V(t)$ is a proxy process for $X(t)$, and $w(t)$ is a weight process.

Discrete-time observation: Existing Methods and Targeting Algorithm

Naive Acceptance-Rejection Algorithm: Simulate process $\tilde{Z}(t)$. Accept the samples that match with the observation (set $W = 1$), otherwise reject ($W = 0$).

$$\hat{\pi}(t,z) = \frac{\sum_{i=1}^{N_s} 1_{\{z\}}(\tilde{Z}^{(i)}(t))\, W^{(i)}}{\sum_{i=1}^{N_s} W^{(i)}} \tag{5}$$

Problem: Most of the weights will be zero.

Conditional Propensity: seeks to sample directly from the conditional distribution via the use of approximated conditional propensity functions.

Targeting Algorithm: We construct a proxy process $\tilde{Z}$ always reaching the target state. Our approach takes into account that the reaction counts between two observation snapshots satisfy linear constraints, and it uses inhomogeneous Poisson processes as proposals for the reaction counts to facilitate exact interpolation.

Estimating Hidden States

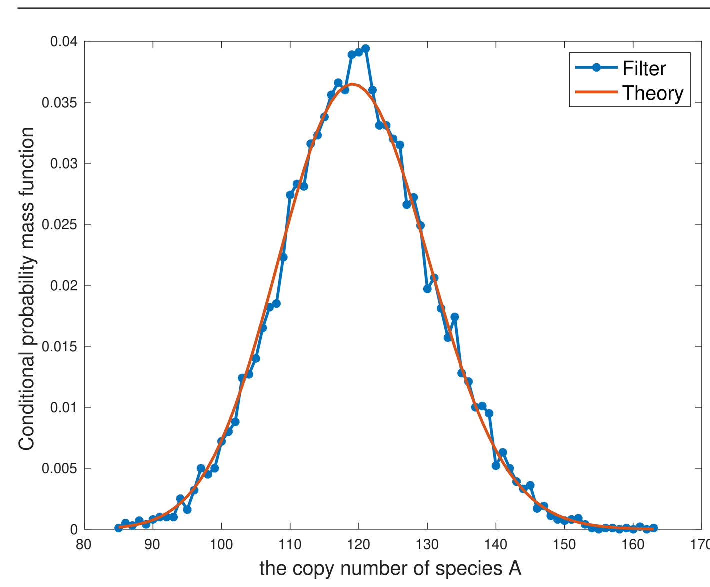

Figure 1. Example 1 (Species S is observed in continuous-time, while the species A is unobserved). Thus, $Z(t) = (X(t), Y(t)) = (\#A(t), \#S(t))$. Estimated conditional probability distribution of $A(T)$ compared with the exact theoretical function at time $T = 20$. $z_0 = (0, 5)$, parameter vector $c = (1, 5, 1)$, filter sample size $N_s = 10,000$. Theoretical distribution is computed precisely by $\pi(t,x) = \frac{\lambda^x e^{-\lambda}}{x!}$, where $\lambda = \int_0^t c_1 y(t) dt$.

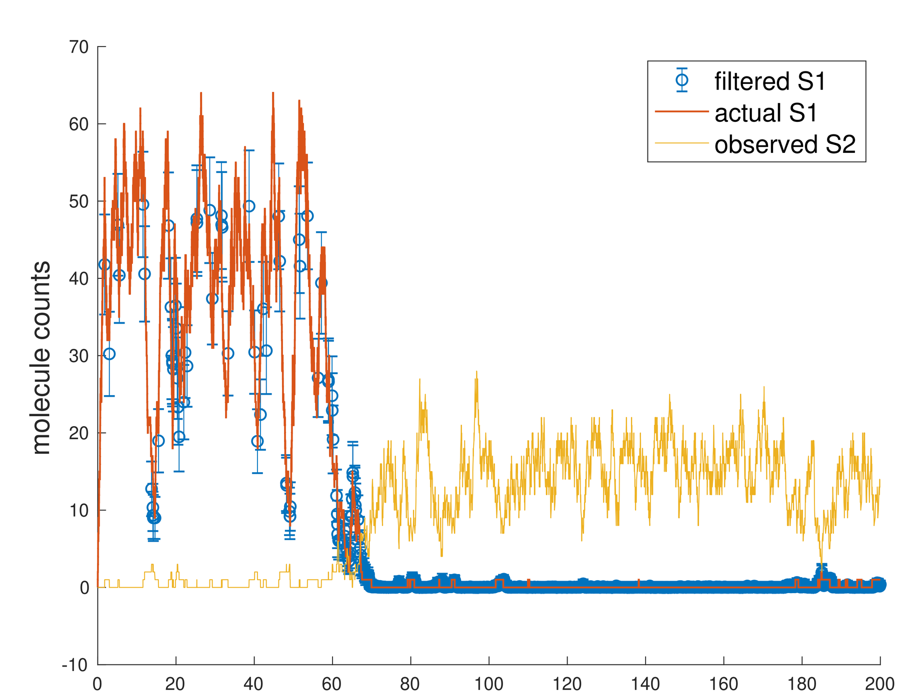

Figure 2. Example 2 (Species S_2 is observed in continuous-time, while the species S_1 is unobserved). Estimation of copy number of S_1 at time t based on one observed trajectory of S_2 over $[0,t]$. Estimates are shown at jump times of the observed species S_2 along with their 68% confidence intervals. Filter sample size $N_s = 10,000$. $Z(0) = (0, 0)$, $\alpha_1 = 50$, $\alpha_2 = 16$, $\beta = 2.5$, $\gamma = 1$.

Bayesian Parameter Estimation

Propensity functions $a_j = a_j(x, y, c)$ could contain unknown constants.

- Start with a prior distribution $\bar{\mu}$ on the parameter space
- Treat c as a constant-valued stochastic process $C(t)$
- Apply our algorithm to Markov process $Z(t) = (X(t), C(t), Y(t))$

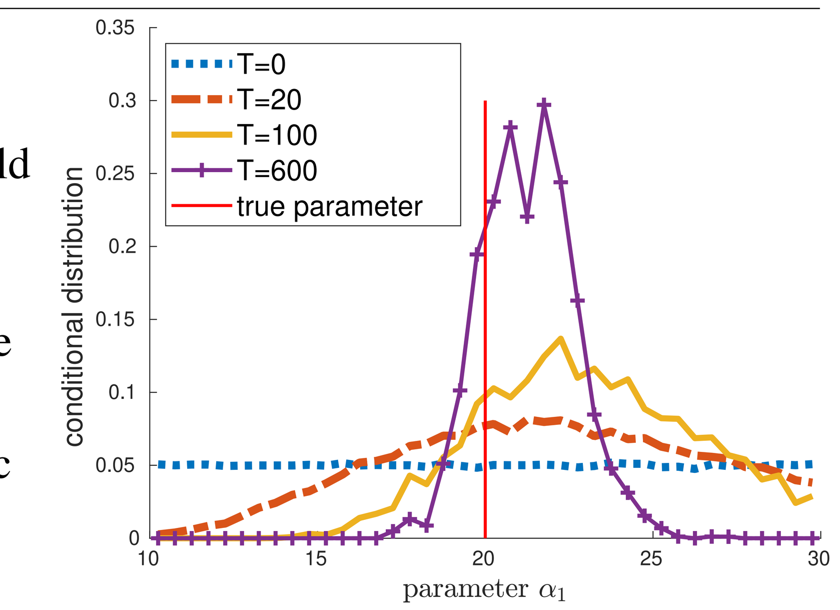

Figure 3. Example 2. The posterior conditional probability density function of parameters α_1 and in the genetic toggle example. The nominal parameter values were $\alpha_1 = 20, \alpha_2 = 9, \beta = 2.5, \gamma = 1$.

Algorithm Performance

We compare our Targeting Algorithm with existing methods.

Method	Observation	Total Variation Error (95% interval)	Effective Sample Fraction (ESF)	Runtime (seconds)
Naive	$x_T = 368$	1.1353, [1.1077, 1.1630]	0.027	2.3
Targeting ($\bar{\lambda}_1$)	(common observation)	0.2037, [0.1995, 0.2080]	0.919	2.6
Targeting ($\bar{\lambda}_2$)		0.2072, [0.2029, 0.2116]	0.900	2.7
CP (approx)		0.3485, [0.3414, 0.3556]	0.322	2.4
CP (exact)		0.1957, [0.1920, 0.1994]	1.000	2.4
Naive	$x_T = 404$	1.9036, [1.8934, 1.9139]	0.002	2.4
Targeting ($\bar{\lambda}_1$)	(rare observation)	0.1979, [0.1942, 0.2016]	0.923	2.6
Targeting ($\bar{\lambda}_2$)		0.2297, [0.2248, 0.2347]	0.654	2.4
CP (approx)		0.3351, [0.3275, 0.3427]	0.324	2.1
CP (exact)		0.1909, [0.1872, 0.1947]	1.000	2.3

Table 1. Example 3. We took $x_0 = 1000$, $c = 2$, $T = 0.5$, $t = 0.2$, $\Delta t = 0.02$. We used $N_r = 100$ trials, while each trial has filter sample size $N_s = 1000$. We note that in our experiment for the rare observation case, the naive method has only 80% of the trials provided meaningful results, and the rest of 20% runs of the naive filter failed when none of the sample landed at the target state, resulting a zero denominator issue. Total Variation Error (TVE) = $\mathbb{E}\left[\sum_z |\hat{\pi}(t,z) - \pi(t,z)| \,\middle|\, Y(T) = y\right]$.

Eeffective Sample Fraction = $\frac{N_e}{N_s} = \frac{1}{N_s}\frac{\left(\sum_{i=1}^{N_s} W^i\right)^2}{\sum_{i=1}^{N_s}(W^i)^2}$.

Highlights

- New particle filtering algorithms for full state estimation based on exact partial state observations of reaction networks.
- Rigorous proof of the algorithms.
- Numerical examples demonstrate the efficacy of the algorithms.

References

[1] Fulvia Confortola and Marco Fuhrman. Filtering of continuous-time markov chains with noise-free observation and applications. *Stochastics An International Journal of Probability and Stochastic Processes*, 85(2):216–251, 2013.

[2] Bert Fristedt, Naresh Jain, and NikolaĭVladimirovich Krylov. *Filtering and Prediction: A Primer:*, volume 38. American Mathematical Soc., 2007.

[3] Andrew Golightly and Chris Sherlock. Efficient sampling of conditioned markov jump processes. *Statistics and Computing*, 29(5):1149–1163, 2019.

[4] Muruhan Rathinam and Mingkai Yu. State and parameter estimation from exact partial state observation in stochastic reaction networks. *The Journal of Chemical Physics*, 154(3), 2021.

[5] Muruhan Rathinam and Mingkai Yu. Stochastic filtering of reaction networks partially observed in time snapshots. *Journal of Computational Physics*, 515:113265, 2024.