Exact MCMC for SDEs with unit diffusion

Xinyi Pei, Minhyeok Kim and **Vinayak Rao**

Department of Statistics, Purdue University

ICERM Workshop *'NPBayesian Inference - Computational Issues'*, Providence, RI

Stochastic differential equations (SDEs)

SDEs: Nonparametric priors over continuous-time paths $X_t, t \in [0, T]$

$$\mathrm{d}X_t = \alpha_\theta(X_t)\mathrm{d}t + \beta_\theta(X_t)\mathrm{d}W_t$$

- X_t: state at time t
- W_t: Brownian motion at time t

Used in finance, physics, biology, chemistry

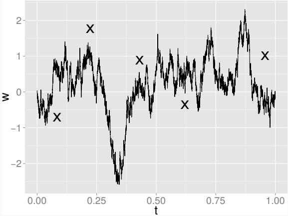

Allow modelers to capture

- mechanistic knowledge about system dynamics
- posterior state/path uncertainty given noisy measurements

Stochastic differential equations (SDEs)

SDEs: Nonparametric priors over continuous-time paths $X_t, t \in [0, T]$

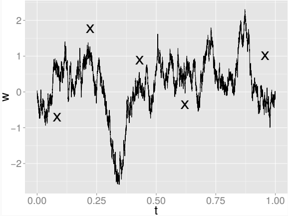

$$\mathrm{d}X_t = \alpha_\theta(X_t)\mathrm{d}t + \beta_\theta(X_t)\mathrm{d}W_t$$

- X_t: state at time t
- W_t: Brownian motion at time t

Used in finance, physics, biology, chemistry

Allow modelers to capture

- mechanistic knowledge about system dynamics
- posterior state/path uncertainty given noisy measurements

Problem: Simulating from these models is very challenging

- Typically, $p(X_{t+\Delta}|X_t)$ is intractable to simulate/evaluate
- Posterior simulation (given noisy observations) is even harder

Discrete-time approximations

Common approaches use **discrete-time approximations**.

E.g. Euler-Maruyama uses the Gaussian approximation
$p(X_{t+\Delta}|X_t) \approx N\big(\alpha_\theta(X_t)\Delta, \beta_\theta(X_t)^2\Delta\big)$ on the grid $[0, \Delta, 2\Delta, \ldots, T]$

- Allow access to discrete-time computational tools

On the downside,

- Solutions are biased from the discrete-time approximation
- Controlling this requires a fine grid and expensive computations
- Simulating on a very fine discretization grid can also result in MCMC mixing issues [Roberts and Stramer, 2001]

(But see *multilevel Monte Carlo* based approaches such as
Rhee, C. H., & Glynn, P. W. (2015). *Unbiased estimation with square root convergence for SDE models.* Operations Research)

This work

We allow **exact** MCMC simulation (with no time-discretization error)

We allow use of standard MCMC tools (from e.g. Gaussian processes)

We consider SDEs with **constant diffusion**:

$$\mathrm{d}X_t = \alpha_\theta(X_t)\mathrm{d}t + \mathrm{d}W_t, \quad \alpha_\theta(x) = \nabla_x A(x)$$

This work

We allow **exact** MCMC simulation (with no time-discretization error)

We allow use of standard MCMC tools (from e.g. Gaussian processes)

We consider SDEs with **constant diffusion**:

$$\mathrm{d}X_t = \alpha_\theta(X_t)\mathrm{d}t + \mathrm{d}W_t, \quad \alpha_\theta(x) = \nabla_x A(x)$$

Assuming constant diffusion entails no real loss of generality in 1-D because of the *Lamperti transform*

- More restrictive in higher dimensions

Nevertheless, this represents a broad and useful class of SDEs.

Exact simulation [Beskos & Roberts, 2005 and follow-up work]

Write $\mathcal{Q}$ and $\mathcal{W}$ for the probability law over paths associated with the SDE of interest and Brownian motion.

From Girsanov's theorem, for any path $X := \{X_t, t \in [0, T]\}$:

$$\frac{\mathrm{d}\mathcal{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}, \quad \phi(\cdot) \geq 0 \tag{1}$$

$$\phi(\cdot) = \frac{1}{2}\left(\alpha^2(\cdot) + \alpha'(\cdot)\right) + \text{const}$$

Exact simulation [Beskos & Roberts, 2005 and follow-up work]

Write $\mathcal{Q}$ and $\mathcal{W}$ for the probability law over paths associated with the SDE of interest and Brownian motion.

From Girsanov's theorem, for any path $X := \{X_t, t \in [0, T]\}$:

$$\frac{\mathrm{d}\mathcal{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}, \quad \phi(\cdot) \geq 0 \tag{1}$$

$$\phi(\cdot) = \frac{1}{2}\left(\alpha^2(\cdot) + \alpha'(\cdot)\right) + \text{const}$$

Really, $\mathcal{W}$ is the law of *biased Brownian bridge* whose end-value X_T is distributed as $p(X_T|X_0) \propto \exp(A(X_T) - \frac{1}{2}(X_T - X_0)^2)$

- We can happily ignore this

Exact simulation [Beskos & Roberts, 2005 and follow-up work]

Write $\mathcal{Q}$ and $\mathcal{W}$ for the probability law over paths associated with the SDE of interest and Brownian motion.

From Girsanov's theorem, for any path $X := \{X_t, t \in [0, T]\}$:

$$\frac{\mathrm{d}\mathcal{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}, \quad \phi(\cdot) \geq 0 \tag{1}$$

$$\phi(\cdot) = \frac{1}{2}\left(\alpha^2(\cdot) + \alpha'(\cdot)\right) + \text{const}$$

> Really, $\mathcal{W}$ is the law of *biased Brownian bridge* whose end-value X_T is distributed as $p(X_T|X_0) \propto \exp(A(X_T) - \frac{1}{2}(X_T - X_0)^2)$
>
> - We can happily ignore this

Eq (1) suggests *rejection sampling* with $\mathcal{W}$ as *proposal distribution*

Idealized simulation of SDE using $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}$

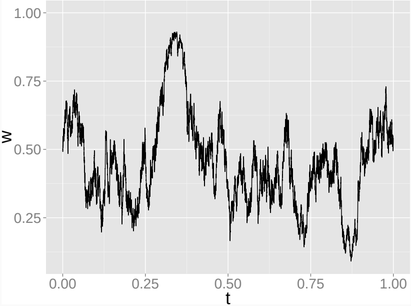

Propose X from the biased Brownian bridge $\mathcal{W}$

Idealized simulation of SDE using $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}$

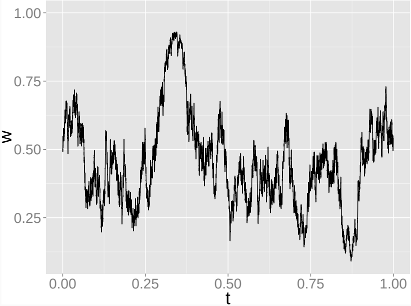

Propose X from the biased Brownian bridge $\mathcal{W}$

Set $\mathsf{acc}(X) = \exp\left\{-\int_0^t \phi(X_t)\mathrm{d}t\right\}$

Accept X with prob. $\mathsf{acc}(X)$

Idealized simulation of SDE using $\frac{\mathrm{d}\mathcal{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}$

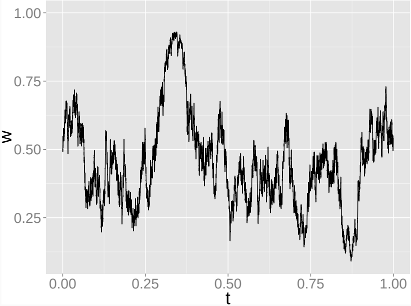

Propose X from the biased Brownian bridge $\mathcal{W}$

Set $\mathtt{acc}(X) = \exp\left\{-\int_0^t \phi(X_t)\mathrm{d}t\right\}$	Simulate a $\phi(X_t)$-intensity Poisson process
Accept X with prob. $\mathtt{acc}(X)$	Accept if zero Poisson events on $[0, T]$

Idealized simulation of SDE using $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathcal{W}}(X) \propto \exp\left\{-\int_0^T \phi(X_t)\mathrm{d}t\right\}$

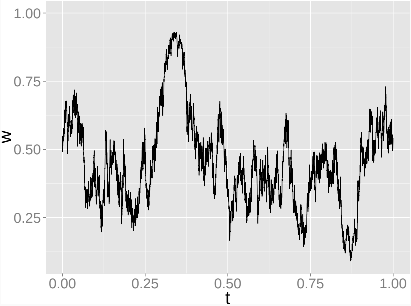

Propose X from the biased Brownian bridge $\mathcal{W}$

Set $\texttt{acc}(X) = \exp\left\{-\int_0^t \phi(X_t)\mathrm{d}t\right\}$	Simulate a $\phi(X_t)$-intensity Poisson process
Accept X with prob. $\texttt{acc}(X)$	Accept if zero Poisson events on $[0, T]$

Remarkably, approach 2 can be done exactly with finite computation!

Exactly simulating a rate-$\phi(X_t)$ Poisson process by thinning

Assume a uniform bound M on $\phi(\cdot)$

- This is only true for so-called EA1 diffusions, but we will relax this

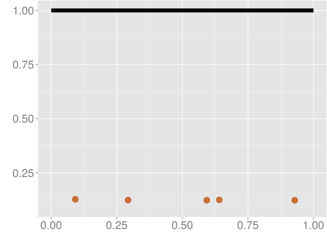

Simulate ψ from a Poisson process with intensity $M \geq \phi(\cdot)$

Exactly simulating a rate-$\phi(X_t)$ Poisson process by thinning

Assume a uniform bound M on $\phi(\cdot)$

- This is only true for so-called EA1 diffusions, but we will relax this

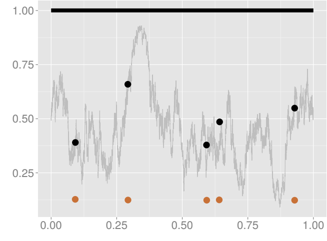

Simulate ψ from a Poisson process with intensity $M \geq \phi(\cdot)$

Evaluate (simulate) X at times ψ from $\mathcal{W}$. Call this X_ψ

Exactly simulating a rate-$\phi(X_t)$ Poisson process by thinning

Assume a uniform bound M on $\phi(\cdot)$

- This is only true for so-called EA1 diffusions, but we will relax this

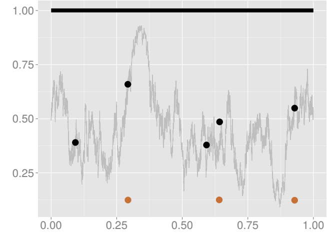

Simulate ψ from a Poisson process with intensity $M \geq \phi(\cdot)$

Evaluate (simulate) X at times ψ from $\mathcal{W}$. Call this X_ψ

Keep time $e \in \psi$ w.p. $\frac{\phi(X_e)}{M}$. *Survivors form a rate-$\phi(X_t)$ Poiss. proc.*

Exactly simulating a rate-$\phi(X_t)$ Poisson process by thinning

Assume a uniform bound M on $\phi(\cdot)$

- This is only true for so-called EA1 diffusions, but we will relax this

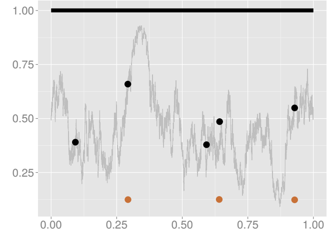

Simulate ψ from a Poisson process with intensity $M \geq \phi(\cdot)$

Evaluate (simulate) X at times ψ from $\mathcal{W}$. Call this X_ψ

Keep time $e \in \psi$ w.p. $\frac{\phi(X_e)}{M}$. *Survivors form a rate-$\phi(X_t)$ Poiss. proc.*

Accept X if no survivors.

- X can be evaluated at other times from Brownian bridge $X \mid X_\psi$

Why Markov chain Monte Carlo (MCMC) alternatives?

This clever idea of Beskos and collaborators can still be inefficient

- Long intervals and strong deviations from Brownian motion can have high rejection rates
- Does not handle posterior simulation (i.e. simulation conditional on observations) naturally

Extended to SDEs with unbounded $\phi(\cdot)$, but these **EA2** and **EA3** (especially) algorithms are computationally challenging (more later)

Why Markov chain Monte Carlo (MCMC) alternatives?

This clever idea of Beskos and collaborators can still be inefficient

- Long intervals and strong deviations from Brownian motion can have high rejection rates
- Does not handle posterior simulation (i.e. simulation conditional on observations) naturally

Extended to SDEs with unbounded $\phi(\cdot)$, but these **EA2** and **EA3** (especially) algorithms are computationally challenging (more later)

Our contribution: an exact MCMC sampler that addresses these issues by producing *dependent* instead of independent samples

Given observations Y of X at times S, let $\ell_Y(X)$ be the likelihood

For any path realization X, let $M(X) \geq \phi(X_t), t \in [0, T]$

Let $\mathcal{M}_{M(X)}$ be the law of a rate-$M(X)$ Poiss. proc. on $[0, T]$

Theorem: *Define a probability $\mathcal{P}$ on the path $\times$ point process space*

$$\mathcal{P}(dX, d\psi) \propto \underbrace{\mathcal{W}(dX)}_{\text{Brownian motion}} \underbrace{\ell_Y(X)}_{\text{Likelihood}} \underbrace{\mathcal{M}_{M(X)}(d\psi)}_{\text{Poisson process}} \underbrace{\prod_{g \in \psi} \left(1 - \frac{\phi(X_g)}{M(X)}\right)}_{\text{Thinning}}.$$

Then the SDE posterior $\mathcal{Q}^Y(dX) \propto \mathcal{Q}(dX)\ell_Y(X)$ is marginal of $\mathcal{P}$.

A Gibbs sampling approach

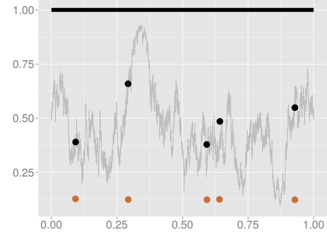

Our approach: A Gibbs sampler targeting $\mathcal{P}(dX, d\psi)$ that alternately:

- Samples a new path X given the point process ψ
- Samples a new point process realization ψ given the path X

A Gibbs sampling approach

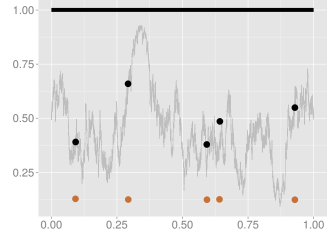

Our approach: A Gibbs sampler targeting $\mathcal{P}(dX, d\psi)$ that alternately:

- Samples a new path X given the point process ψ
- Samples a new point process realization ψ given the path X

MCMC state space consists of ψ, $X_{\psi\cup S}$ and $M(X)$

- As before, the SDE path X can be simulated anywhere else as Brownian motion conditioned on $(X_{\psi\cup S}, M(X))$

Gibbs step: Updating $\psi \mid X$

Proposition 1: *Given trajectory X and observations Y, the random locations ψ are a Poisson process with intensity $M(X) - \phi(X_t)$*

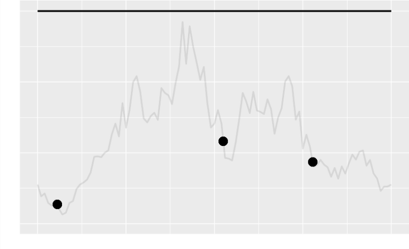

Easy to simulate by Poisson thinning

Note: our MCMC state space only has $(M(X), X_{\psi \cup S})$

Gibbs step: Updating $\psi \mid X$

Proposition 1: *Given trajectory X and observations Y, the random locations ψ are a Poisson process with intensity $M(X) - \phi(X_t)$*

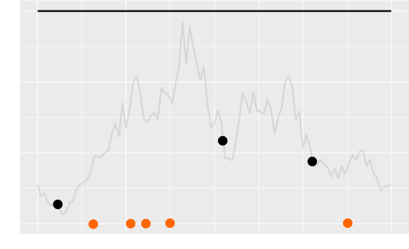

Easy to simulate by Poisson thinning

Note: our MCMC state space only has $(M(X), X_{\psi \cup S})$

- Simulate a Poisson process with intensity $M(X)$

Proposition 1: *Given trajectory X and observations Y, the random locations ψ are a Poisson process with intensity $M(X) - \phi(X_t)$*

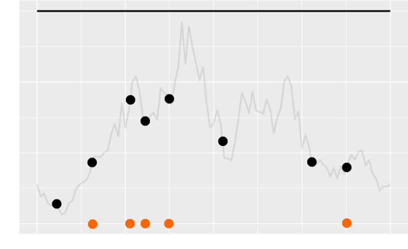

Easy to simulate by Poisson thinning

Note: our MCMC state space only has $(M(X), X_{\psi \cup S})$

- Simulate a Poisson process with intensity $M(X)$
- Conditionally simulate Brownian motion on these times

Proposition 1: *Given trajectory X and observations Y, the random locations ψ are a Poisson process with intensity $M(X) - \phi(X_t)$*

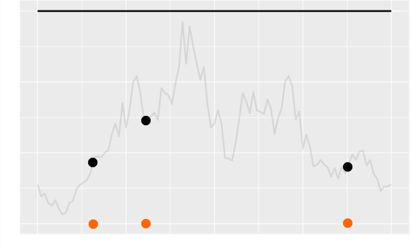

Easy to simulate by Poisson thinning

Note: our MCMC state space only has $(M(X), X_{\psi \cup S})$

- Simulate a Poisson process with intensity $M(X)$
- Conditionally simulate Brownian motion on these times
- Keep Poisson events with probability $1 - \phi(X_t)/M(X)$

Gibbs iterations

Proposition 2: *Conditioned on locations ψ and observations Y, the SDE path X follows a Gaussian process posterior corresponding to:*

- *A biased Brownian bridge prior $\mathcal{W}$ on X*
- *Likelihood that includes observations Y and Poisson events ψ*

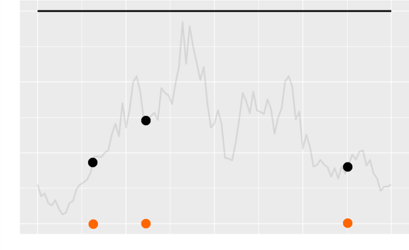

Can use standard MCMC tools like Hamiltonian Monte Carlo, elliptical slice sampling, data augmentation for logit/probit models

- We only need to update $X_{\psi \cup S}, M(X)$

Gibbs iterations

Proposition 2: *Conditioned on locations ψ and observations Y, the SDE path X follows a Gaussian process posterior corresponding to:*

- *A biased Brownian bridge prior $\mathcal{W}$ on X*
- *Likelihood that includes observations Y and Poisson events ψ*

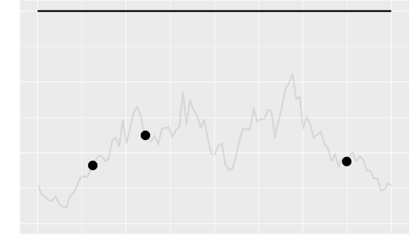

Can use standard MCMC tools like Hamiltonian Monte Carlo, elliptical slice sampling, data augmentation for logit/probit models

- We only need to update $X_{\psi \cup S}, M(X)$

Now we must include a **path-dependent bound** $M(X)$ on $\phi(\cdot)$

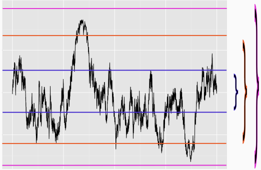

Given fixed, nested intervals $\mathcal{I} = \{(L_i, U_i)\}$, the **layer** $X^{\Updownarrow}$ of a Brownian path X on $[0, T]$ is the smallest interval containing it

- Simulating $X^{\Updownarrow}$ for Brownian X isn't hard

What about general SDEs where $\phi(\cdot)$ is unbounded?

Now we must include a **path-dependent bound** $M(X)$ on $\phi(\cdot)$

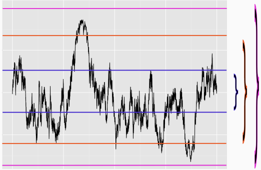

Given fixed, nested intervals $\mathcal{I} = \{(L_i, U_i)\}$, the **layer** $X^{\Updownarrow}$ of a Brownian path X on $[0, T]$ is the smallest interval containing it

- Simulating $X^{\Updownarrow}$ for Brownian X isn't hard

Set $M(X) = M(X^{\Updownarrow}) \overset{\text{def}}{\geq} \sup_{x \in [L_{X^{\Updownarrow}}, U_{X^{\Updownarrow}}]} \phi_\theta(x) \geq \sup_{t \in [0,T]} \phi_\theta(X_t)$.

What about general SDEs where $\phi(\cdot)$ is unbounded?

Now we must include a **path-dependent bound** $M(X)$ on $\phi(\cdot)$

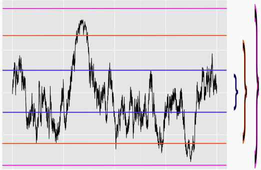

Given fixed, nested intervals $\mathcal{I} = \{(L_i, U_i)\}$, the **layer** $X^{\Updownarrow}$ of a Brownian path X on $[0, T]$ is the smallest interval containing it

- Simulating $X^{\Updownarrow}$ for Brownian X isn't hard

Set $M(X) = M(X^{\Updownarrow}) \overset{\text{def}}{\geq} \sup_{x \in [L_{X^{\Updownarrow}}, U_{X^{\Updownarrow}}]} \phi_\theta(x) \geq \sup_{t \in [0,T]} \phi_\theta(X_t)$.

Now we can simulate rate-$\phi(X_t)$ Poisson process by thinning.

The exact (EA3) rejection sampler for the general case:

1. Simulate the layer $X^{\Updownarrow}$ of the Brownian path X
2. Simulate ψ from a rate-$M(X^{\Updownarrow})$ Poisson process
3. Simulate $X_\psi | X^{\Updownarrow}$
4. Keep each event $e \in \psi$ with probability $\phi(X_e)/M(X^{\Updownarrow})$
5. Accept X if all events in ψ are thinned, else restart

What about general SDEs where $\phi(\cdot)$ is unbounded?

Now we must include a **path-dependent bound** $M(X)$ on $\phi(\cdot)$

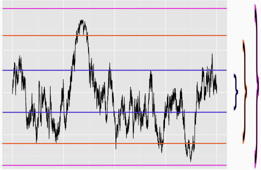

Given fixed, nested intervals $\mathcal{I} = \{(L_i, U_i)\}$, the **layer** $X^{\Updownarrow}$ of a Brownian path X on $[0, T]$ is the smallest interval containing it

- Simulating $X^{\Updownarrow}$ for Brownian X isn't hard

Set $M(X) = M(X^{\Updownarrow}) \overset{\text{def}}{\geq} \sup_{x \in [L_{X^{\Updownarrow}}, U_{X^{\Updownarrow}}]} \phi_\theta(x) \geq \sup_{t \in [0,T]} \phi_\theta(X_t)$.

Now we can simulate rate-$\phi(X_t)$ Poisson process by thinning.

The exact (EA3) rejection sampler for the general case:

1. Simulate the layer $X^{\Updownarrow}$ of the Brownian path X
2. Simulate ψ from a rate-$M(X^{\Updownarrow})$ Poisson process
3. Simulate $X_\psi | X^{\Updownarrow}$ (this is very challenging!)
4. Keep each event $e \in \psi$ with probability $\phi(X_e)/M(X^{\Updownarrow})$
5. Accept X if all events in ψ are thinned, else restart

Benefits of our proposed Gibbs approach

Our MCMC approach lets us reorder computations to avoid this

- Recall, our Gibbs step updates X *conditioned* on ψ

The EA3 rejection smplr must simulate $X^{\Updownarrow}$ **first** to simulate ψ and X_ψ

Our Gibbs step can simulate $X_{\psi \cup S}$ and ***then*** $X^{\Updownarrow}$

- Much easier, esp. in high-freq. settings, where X_S is high-dim

Benefits of our proposed Gibbs approach

Our MCMC approach lets us reorder computations to avoid this

- Recall, our Gibbs step updates X *conditioned* on ψ

The EA3 rejection smplr must simulate $X^{\Updownarrow}$ **first** to simulate ψ and X_ψ

Our Gibbs step can simulate $X_{\psi \cup S}$ and ***then*** $X^{\Updownarrow}$

- Much easier, esp. in high-freq. settings, where X_S is high-dim

We show updating $X \mid \psi$ involves the following steps:

- Update $\widetilde{X}_S \mid X_S$ from a Markov kernel targeting $p_{Z_h}(X_S) l_Y(X_S)$ (easy with Gaussian noise, else use e.g. HMC, Polya-Gamma)
- Simulate $\widetilde{X}_\psi \mid \widetilde{X}_S$ from a Brownian bridge
- Simulate the layer $\widetilde{X}^{\Updownarrow}$ of a Brownian bridge given $\widetilde{X}_{\psi \cup S}$
- Accept/reject with probability Poiss($\psi \mid M(X^{\Updownarrow})$)

EA3 example (double-well potential Langevin process)

SDE: $$dX_t = (-pX_t^3 + qX_t)dt + dW_t$$

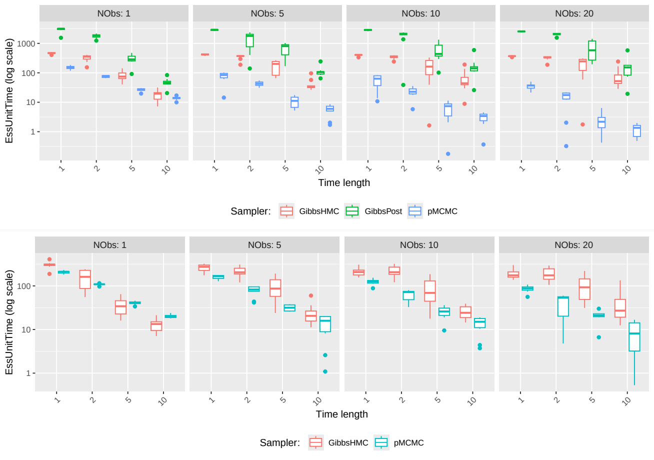

(Effective samples per second) vs (Interval length T) and (Number of observations $|Y|$) under Gaussian (top) and Poisson (bottom) noise

Conclusions

Described a framework for exact MCMC over a useful class of SDEs

Not discussed here:

- *Additional* Poisson process ξ to make the $\psi|X$ Gibbs update easy
- Novel extensions to parameter inference
- Empirical evaluations, including a dataset of oxygen isotopic levels in ice ($\delta^{18}O$) in Greenland over the last 60,000 years

Our paper should be up on `arxiv` in a week.

Meanwhile, for the EA1 setting (with bounded $\phi(\cdot)$), you can look at:

- Wang, **Rao**, Teh, *An exact auxiliary variable Gibbs sampler for a class of diffusions*, J. Comput. Graph. Stat., 2020

Conclusions

Described a framework for exact MCMC over a useful class of SDEs

Not discussed here:

- *Additional* Poisson process ξ to make the $\psi|X$ Gibbs update easy
- Novel extensions to parameter inference
- Empirical evaluations, including a dataset of oxygen isotopic levels in ice ($\delta^{18}O$) in Greenland over the last 60,000 years

Our paper should be up on `arxiv` in a week.

Meanwhile, for the EA1 setting (with bounded $\phi(\cdot)$), you can look at:

- Wang, **Rao**, Teh, *An exact auxiliary variable Gibbs sampler for a class of diffusions*, J. Comput. Graph. Stat., 2020

Thank you!

Thank you!

Exact simulation [Beskos & Roberts, 2005]

Write $\mathbb{W}$ and $\mathbb{Q}$ for the probability law over paths associated with the Brownian motion (really a biased Brownian bridge) and the SDE.

From Girsanov's theorem (with $A(u) = \int_0^u \alpha(t)\mathrm{d}t$):

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{W}}(\omega) = \exp\left\{A(\omega_T) - A(x) - \frac{1}{2}\int_0^T \left(\alpha^2(\omega_t) + \alpha'(\omega_t)\right)\mathrm{d}t\right\}$$

Under fairly mild assumptions, we can write this as

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{W}}(\omega) \propto \exp\left\{-\int_0^T \phi(\omega_t)\mathrm{d}t\right\}, \quad \phi(\cdot) \geq 0$$

- $\phi(\cdot)$ is a function of the drift $\alpha(\cdot)$, and is usually easy to evaluate

Write $\mathbb{W}$ and $\mathbb{Q}$ for the probability law over paths associated with the Brownian motion (really a biased Brownian bridge) and the SDE.

From Girsanov's theorem (with $A(u) = \int_0^u \alpha(t)\mathrm{d}t$):

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{W}}(\omega) = \exp\left\{A(\omega_T) - A(x) - \frac{1}{2}\int_0^T \left(\alpha^2(\omega_t) + \alpha'(\omega_t)\right)\mathrm{d}t\right\}$$

Under fairly mild assumptions, we can write this as

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{W}}(\omega) \propto \exp\left\{-\int_0^T \phi(\omega_t)\mathrm{d}t\right\}, \quad \phi(\cdot) \geq 0$$

- $\phi(\cdot)$ is a function of the drift $\alpha(\cdot)$, and is usually easy to evaluate

Suggests *rejection sampling* with $\mathbb{W}$ as *proposal distribution*

Proposed method with observations

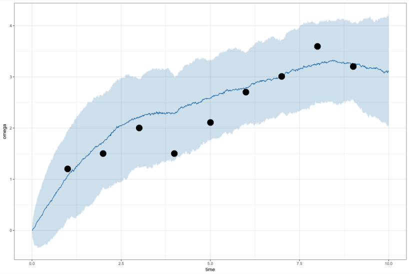

The method of Beskos et al does not easily extend to observations.

We compared instead with an approximate sampler using pMCMC on the time-discretized approximation to the SDE.

What if we don't have a uniform bound on $\phi(\cdot)$?

Rather than $\lim_{|x|\to\infty} \phi(x) < \infty$, we allow $\lim_{x\to-\infty} \phi(x) = \infty$

What if we don't have a uniform bound on $\phi(\cdot)$?

Rather than $\lim_{|x|\to\infty} \phi(x) < \infty$, we allow $\lim_{x\to-\infty} \phi(x) = \infty$

Now our MCMC state-space must include the path minimum

- MCMC state-space: $(\Psi, \omega_\Psi, \omega_{\text{min}}, t_{\text{min}})$

What if we don't have a uniform bound on $\phi(\cdot)$?

Rather than $\lim_{|x|\to\infty} \phi(x) < \infty$, we allow $\lim_{x\to-\infty} \phi(x) = \infty$

Now our MCMC state-space must include the path minimum

- MCMC state-space: $(\Psi, \omega_\Psi, \omega_{\text{min}}, t_{\text{min}})$

Conditioned on its minimum, Brownian motion is a Bessel process

- Now, updating updating ω given Ψ is much more challenging

What if we don't have a uniform bound on $\phi(\cdot)$?

Rather than $\lim_{|x|\to\infty} \phi(x) < \infty$, we allow $\lim_{x\to-\infty} \phi(x) = \infty$

Now our MCMC state-space must include the path minimum

- MCMC state-space: $(\Psi, \omega_\Psi, \omega_{\text{min}}, t_{\text{min}})$

Conditioned on its minimum, Brownian motion is a Bessel process

- Now, updating updating ω given Ψ is much more challenging

In our paper, we propose a novel 'double accept' MCMC algorithm

- Allows the use standard MCMC tools for Gaussian processes at the cost of including an accept/reject step in the Metropolis-Hastings *proposal*