A Gaussian approximation for stochastic non-linear dynamical processes with annihilation

Wim Wiegerinck, Bert Kappen

SNN, Radboud University Nijmegen

The research reported here is part of the Interactive Collaborative Information Systems (ICIS) project, supported by the Dutch Ministry of Economic Affairs, grant BSIK03024.

Contents

- Motivation: stochastic optimal control
- Stochastic process with annihilation
- Relation with Kalman smoothing
- Path integral formulation
- Mode (optimal path), fluctuations + partition function
- Discussion

Stochastic optimal control problem

- Consider a system with controlled stochastic dynamics

$$dx = (b(x,t) + u)dt + d\xi \qquad d\xi \sim N(0, \nu dt)$$

with control u.

- Find the control $u(.)$ that minimizes the *expected cost* to end-time t_f

$$C(x_0, t_0, u(.)) = \left\langle \int_{t_0}^{t_f} \frac{1}{2} u(x(t),t)^2 + V(x(t),t)\, dt \right\rangle$$

  - u^2 control costs
  - V: path costs

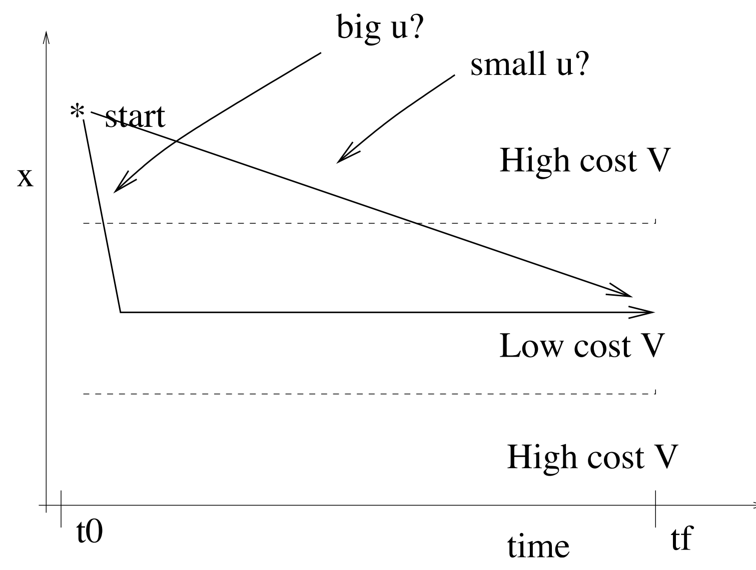

Hamilton-Jacobi-Bellman equation

- Optimal (expected) cost-to-go

$$J(x,t) = \min_{u(.)} C(x,t,u(.)).$$

- J satisfies the HJB eqn.,

$$-\partial_t J = \min_u \left(\frac{1}{2}u^2 + (b+u)\partial_x J + \frac{1}{2}\nu\partial_x^2 J + V \right)$$

with end-condition $J(x,t_f) = 0$.

- The minimization with respect to u yields

$$u = -\partial_x J,$$

$$-\partial_t J = -\frac{1}{2}(\partial_x J)^2 + b\partial_x J + \frac{1}{2}\nu\partial_x^2 J + V$$

Log transformation and optimal control

The non-linear PDE of J can transformed into a linear one by the log transform (W. Flemming, 1978, Kappen 2005). Set

$$J(x,t) = -\nu \log Z(x,t)$$

then the "partition function" Z can be written as

$$Z(x,t) = \int dy \rho(y, t_f | x, t)$$

in which ρ satisfies the linear pde

$$\partial_{t'} \rho(x',t'|x,t) = -\partial_{x'}(b(x',t')\rho(x',t'|x,t)) + \frac{1}{2}\nu\partial^2_{x'}\rho(x',t'|x,t) - \frac{V(x',t')}{\nu}\rho(x',t'|x,t).$$

with begin condition $\rho(x',t|x,t) = \delta(x'-x)$

Fokker-Planck with decay

Goal: compute $\rho(x, t_f|x_0, t_0)$, where

- $\rho(x', t_0|x, t_0) = \delta(x' - x)$
- Evolution according to

$$\partial_t \rho(x, t|x_0, t_0) = -\partial_x(b(x,t)\rho(x,t|x_0,t_0)) + \frac{1}{2}\nu\partial_x^2\rho(x,t|x_0,t_0) - V(x,t)\rho(x,t|x_0,t_0).$$

  - $V = 0 \rightarrow$ reduces to the Fokker-Planck equation, modeling a process of drift and diffusion, due to the terms with $b(x,t)$ and ν respectively.
  - The extra term with the potential V makes that "probability" is not conserved.

A stochastic dynamical process with annihilation

FP with decay describes the following stochastic proces with annihilation: particles start at $x = x_0$ and evolve according

$$dx = b(x,t)dt + d\xi \qquad\qquad d\xi \sim N(0, \nu dt)$$

$$x = x + dx, \quad \text{with probability } 1 - V(x,t)dt$$

$$x = \text{annihilated with probability } V(x,t)dt$$

Example:
$b = 0, V = \frac{1}{2}x^2$
Red : annihilated
Black: survived until t_f

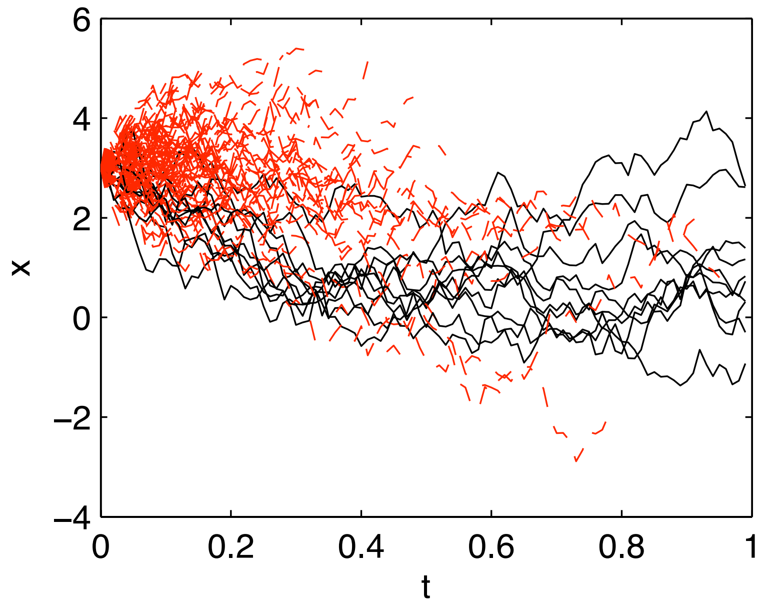

Relation with discrete time Kalman smoothing

Dynamical system equations

$$x_{t+1} = x_t + b(x_t, t) + \epsilon \qquad \epsilon \sim N(0, \nu) \qquad \text{System dynamics}$$
$$y_t = g(x_t) + \eta \qquad \text{Observations}$$

Smoothing

$$p(x_{1:T}|y_{1:T}) \sim \prod_t p(x_{t+1}|x_t)p(y_t|x_t) = \prod_t p(x_{t+1}|x_t)\exp(-V(x_t, t))$$

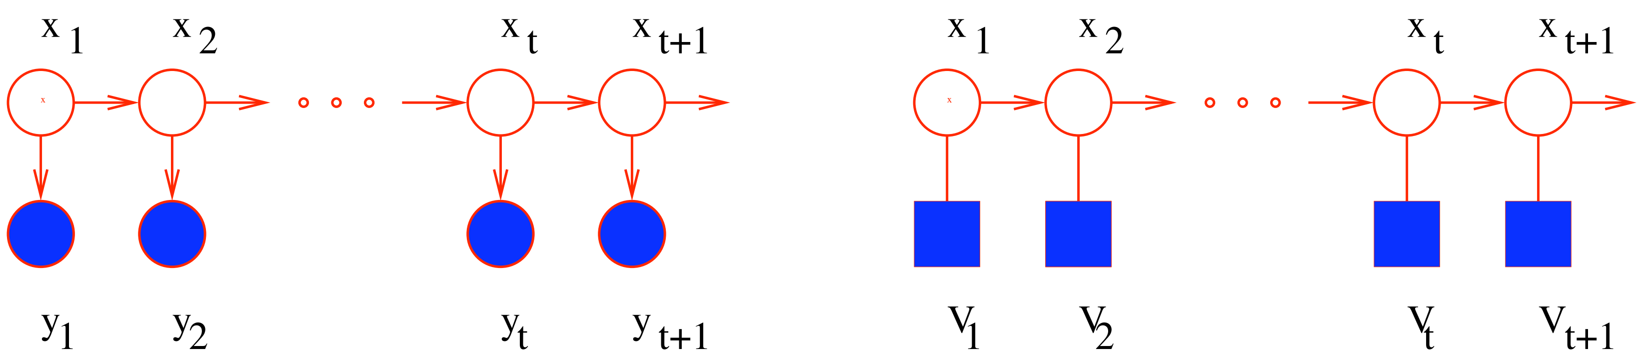

- *Rejection sampling*: sample from dynamics $p(x_{t+1}|x_t)$, reject samples at time t with probability $1 - \exp(-V(x, t))$

The transition density

The transition density state x to y over an infinitesimal time step Δt

$$\rho(y, t+\Delta t|x, t) \propto$$
$$\exp\left(-\left[\frac{(y-x-b(x,t)\Delta t)^2}{2\nu\Delta t} + V(x,t)\Delta t\right]\right)$$

Over n infinitesimal time steps Δt

$$\rho(x_n, t_n|x_0, t_0) \propto$$
$$\int \prod_{i=1}^{n-1} dx_i \exp\left(-\Delta t \sum_{i=0}^{n-1}\left[\frac{1}{2\nu}\left(\frac{x_{i+1}-x_i}{\Delta t} - b(x_i, t_i)\right)^2 + V(x_{i+1}, t_{i+1})\right]\right)$$

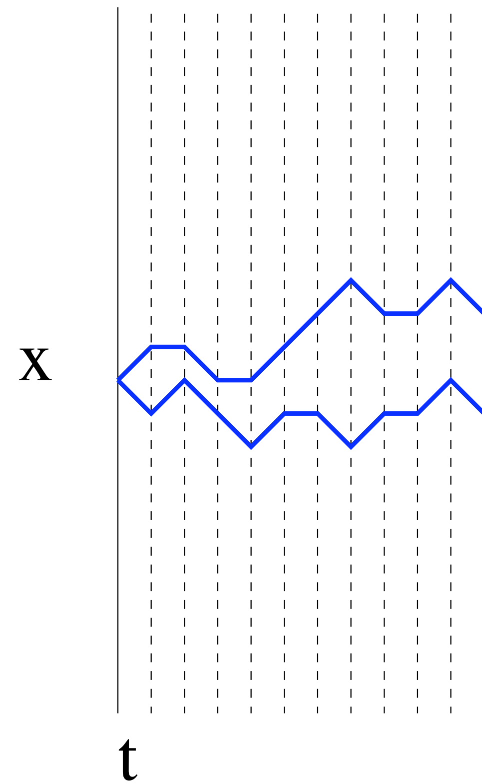

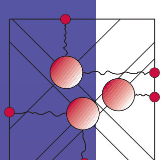

Path integral formulation

In the limit: $\Delta t \sum_{i=0}^{n-1} \to \int_t^{t_f} d\tau$, and $\int \prod_{i=1}^{n-1} dx_i$ becomes an integral over paths that start at x and end at y, denoted as $\int [dx]$.

$$\rho(y, t_f | x_0, t_0) = \int [dx]_x^y \exp\left(-S[x]\right)$$

$$\begin{aligned} S[x] &= \int_{t_0}^{t_f} \left(\frac{\left(\dot{x}(\tau) - b(x(\tau), \tau)\right)^2}{2\nu} + V(x(\tau), \tau) \right) d\tau \\ &= \int_{t_0}^{t_f} L(x(\tau), \dot{x}(\tau), \tau) d\tau \end{aligned}$$

S is called the action, and L the Lagrangian.

Euler-Lagrange equations

The mode of the process. is the path $x(t_0 \to t_f)$, starting at given x_0 and ending at arbitrary y, that minimizes the action S. We do this by applying variational calculus.
Defining "momentum" as

$$p(t) \equiv \partial_{\dot{x}} L(t, x, \dot{x})$$

the optimal path satisfies the well-known Euler-Lagrange equations

$$\begin{aligned} d/dt\, x &= \dot{x} \\ d/dt\, p &= \partial_x L \end{aligned}$$

with begin condition for x (from the problem formulation) and an end-condition an end condition for p (which followed from the variational computation),

$$\begin{aligned} x(t_0) &= x_0 \\ p(t_f) &= 0\ . \end{aligned}$$

Euler-Lagrange equations

In our problem, the Lagrangian is

$$L(x, \dot{x}, t) = \frac{(\dot{x} - b(x,t))^2}{2\nu} + V(x,t)$$

The "momentum" $p(t) \equiv \partial_{\dot{x}} L(t, x, \dot{x}) = \nu^{-1}(\dot{x} - b(x,t))$, then the E-L eqns

$$\begin{aligned} \dot{x}(t) &= b(x,t) + \nu p(t) \\ \dot{p}(t) &= \partial_x V \end{aligned}$$

- Contribution of momentum proportional to noise: Thanks to the fluctuations the surviving particles avoided from running into regions of high annihilation rate and escaped to regions with lower annihilation rate.

A formal forward-backward algorithm

Solution formally found by forward-backward algorithm:

// ** Forward pass ** //
for all initial momenta p_0 **do**
  prepare the system in $(x(t_0) = x_0, p(t_0) = p_0)$
  integrate forwards in time $t_0 \rightarrow t_f$
  if $p(t_f) = 0$ **then**
    keep $x_f = x(t_f)$
  end if
end for
// ** Backward pass ** //
for all kept end states **do**
  prepare the system in $(x(t_f) = x_f; p(t_f) = 0)$
  propagate backwards in time $t_0 \leftarrow t_f$
  return $x_{opt}(t) = x(t)$
end for

Numerical example

We consider a system with $V(x) = \frac{1}{2}\alpha x^2$ and $b(x,t) = 0$. The optimal path can be computed, $x(t) = \frac{\cosh((\alpha\nu)^{1/2}(t_f - t))}{\cosh((\alpha\nu)^{1/2}(t_f))} x_0$

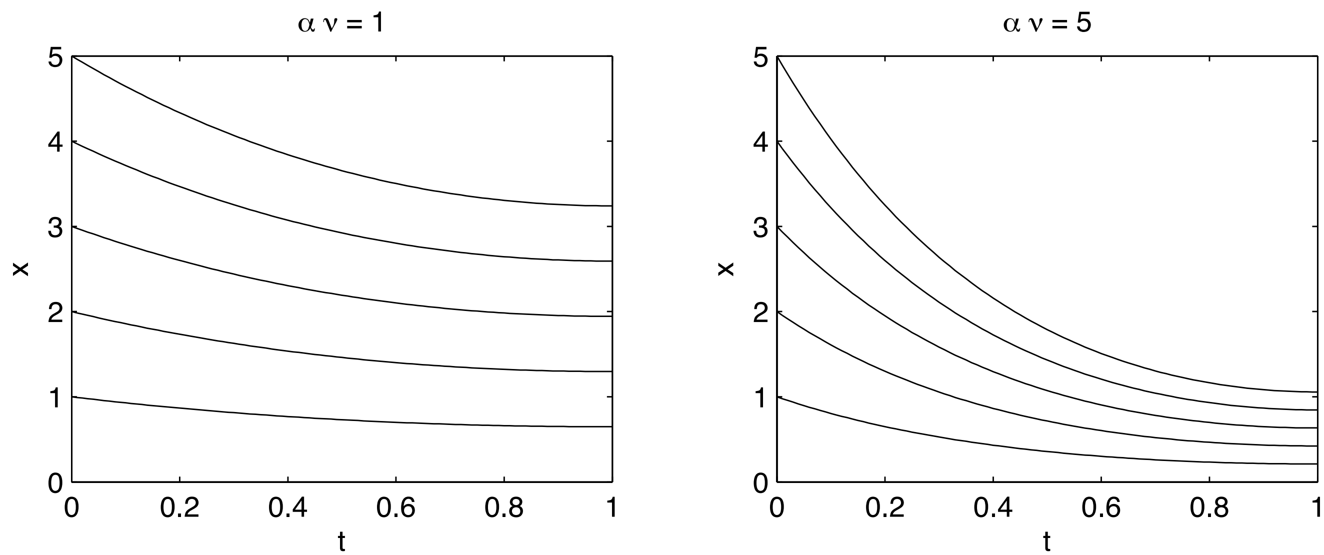

optimal paths starting at different initial points x_0 with $\alpha\nu = 1$ (left) and $\alpha\nu = 5$ (right).

Relation with classical mechanics ($b = 0$)

Stochastic system	Classical mechanics
$p = \nu^{-1}\dot{x}$	$p = m\dot{x}$
$L = \frac{\nu p^2}{2} + V$	$L = \frac{p^2}{2m} - V$
$d/dt\, p = \partial_x V$	$d/dt\, p = -\partial_x V$
$x(t_0) = x_0; \;\; p(t_f) = 0$	$x(t_0) = x_0; \;\; p(t_0) = p_0 \;\;$ (e.g.0)
$H = \frac{\nu p^2}{2} - V$	$H = \frac{p^2}{2m} + V$
Typically, start with large V and large p in direction of $\min V$. End with small V and zero p.	Typically, particles start with large V and zero p. They end with smaller V and larger p, or large V and small p

Size of fluctuations: linear noise approximation

- Fluctuations dominate in short time scale ($d\xi \propto \sqrt{dt}$)
- Drift and annihilation dominate in long time scale ($\propto dt$)
- Drift + state dependent annihilation $\rightarrow$ effective drift described by optimal path + state independent annihilation

$$dx = (b + \nu p)dt + \nu d\xi \equiv \beta(x,t)dt + d\xi \tag{1}$$

Dynamics of fluctuations $\sigma^2(t)$ around mode follows from (1)

$$\partial_t \sigma^2(t) \; = \; 2\,\partial_x \beta(x,t)\,\sigma^2(t) + \nu$$

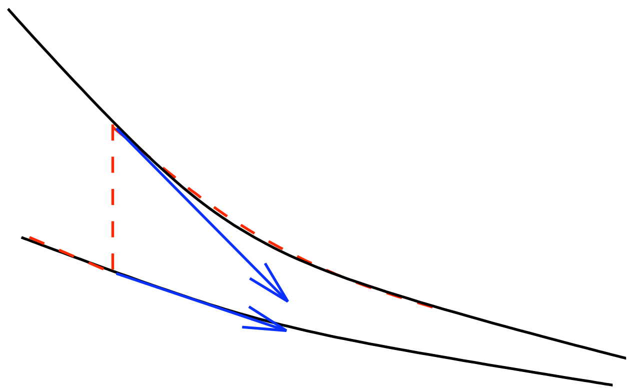

Numerical example

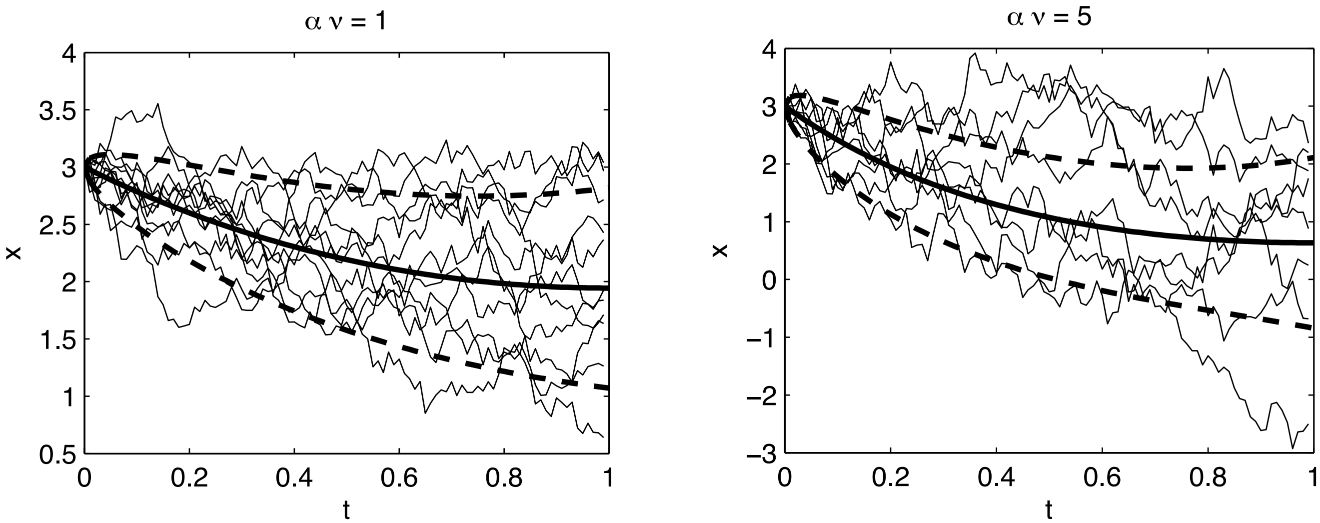

$b = 0$, $V(x) = \frac{1}{2}\alpha x^2$. Optimal paths starting at different initial points x_0 with $\alpha\nu = 1$ (left) and $\alpha\nu = 5$ (right). Bottom: optimal paths (fat lines) starting at $x_0 = 3$, plus indications of esitmated noise $\sigma(t)$ (fat dashed) and some random paths, with $\nu = 1$ (left) and $\nu = 5$ (right). $\alpha = 1$ in both cases. Note: The simulations with $\nu = 1$ started with 500 particles. The one with $\nu = 5$ started with 200 particles.

Partition function

Normalization constant = fraction of particles that survive the process
Approximation

- effective decay rate: fraction of particles that fluctuate towards path $\times$ fraction of particles that survive decay along optimal path

$$V^{\text{path}}(x,t) = \frac{(\beta(x,t) - b(x,t))^2}{2\nu} + V(x,t)\,,$$

- Correction for fluctuations around optimal path

$$V^{\text{path, corrected}}(x,t) = \left\langle \frac{(\beta(x,t) - b(x,t))^2}{2\nu} + V(x,t) \right\rangle_{[x_{\text{opt}},\sigma^2]} \,.$$

Partition function: numerical result

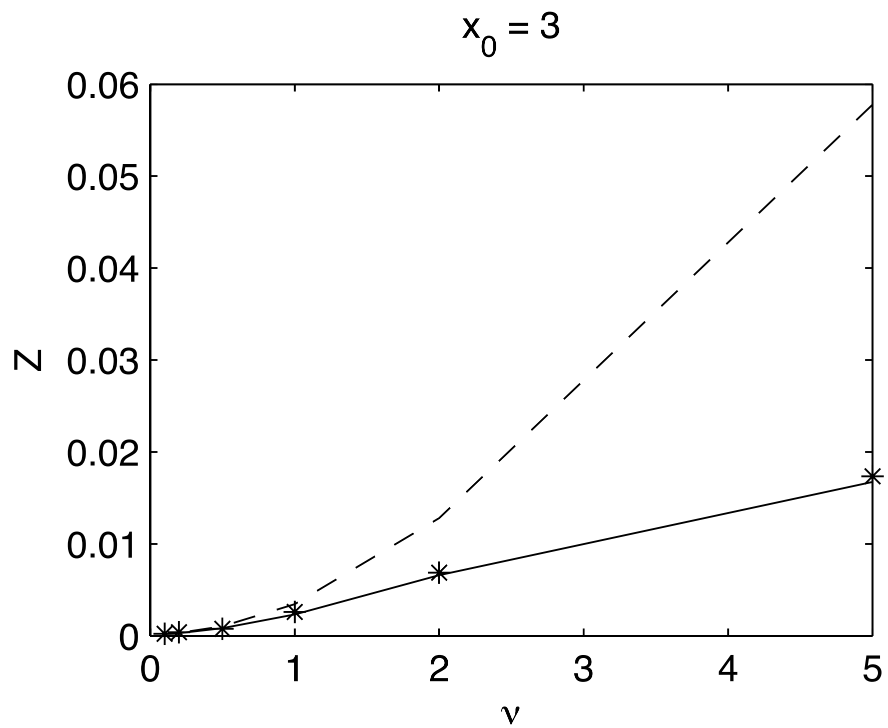

Estimate of the partition function (i.e. fraction of surviving particles) Z based on the mode (dashed) and with Gaussian corrections (drawn) as function of the noise ν. All processes started at $x = 3$. Estimates are compared with results of stochastic simulations, each starting with 100000 particles (stars).

Summary

- Stochastic diffusion with annihilation
- Relevant for:
  - Stochastic optimal control
  - Continuous-time Kalman smoothing (?)
- Path integral formalism
  - Gaussian approximation,
  - mode: optimal path, Euler Lagrange equations
  - fluctuations
  - partition function
- Numerical result for zero drift and quadratic potential

Discussion

- Methods to solve the Euler Lagrange eqns
- Performance on more interesting potentials
- More general stochastic dynamical systems
- Applications of continuous time smoothing