Score matching and nonparametric estimators of drift functions for stochastic differential equations

Manfred Opper

joint work with Philip Batz, Andreas Ruttor,

April 13, 2016

Estimating probability densities by Score Matching

- Let $q_\theta(\cdot)$ be a family of probability densities. Try to estimate $p(x)$ by finding 'best' θ. using samples $x_i \sim p(x)$ for $i = 1, \ldots, n$.

Estimating probability densities by Score Matching

- Let $q_\theta(\cdot)$ be a family of probability densities. Try to estimate $p(x)$ by finding 'best' θ. using samples $x_i \sim p(x)$ for $i = 1, \ldots, n$.
- Maximum likelihood (and Bayes) estimation often suffers from **normalisation** problem $q_\theta(x) = \frac{\tilde{q}_\theta(x)}{Z_\theta}$ with intractable Z_θ.

Estimating probability densities by Score Matching

- Let $q_\theta(\cdot)$ be a family of probability densities. Try to estimate $p(x)$ by finding 'best' θ. using samples $x_i \sim p(x)$ for $i = 1, \ldots, n$.
- Maximum likelihood (and Bayes) estimation often suffers from **normalisation** problem $q_\theta(x) = \frac{\tilde{q}_\theta(x)}{Z_\theta}$ with intractable Z_θ.
- Score matching: the basic identity (Hyvärinen, 2005)

$$J(p\|q_\theta) \doteq \frac{1}{2} \int p(x) \left\|\nabla \ln p(x) - \nabla \ln q_\theta(x)\right\|^2$$

$$= \int p(x) \left\{ \frac{1}{2} \left\|\nabla \ln q_\theta(x)\right\|^2 + \nabla^2 \ln q_\theta(x) \right\} + \text{const}$$

- Use minimisation of empirical loss

$$\sum_{i=1}^{n} \left\{ \frac{1}{2} \left\| \nabla \ln q_\theta(x_i) \right\|^2 + \nabla^2 \ln q_\theta(x_i) \right\}$$

independent of Z_θ !

- Use minimisation of empirical loss

$$\sum_{i=1}^{n} \left\{ \frac{1}{2} \left\| \nabla \ln q_\theta(x_i) \right\|^2 + \nabla^2 \ln q_\theta(x_i) \right\}$$

independent of Z_θ !

- Nonparametric extension (Sriperumbudur, Fukumizu, Kumar, Gretton and Hyvärinen, 2014). Set $\psi(x) \doteq \ln \tilde{q}(x)$

$$\sum_{i=1}^{n} \left\{ \frac{1}{2} \left\| \nabla \psi(x_i) \right\|^2 + \nabla^2 \psi(x_i) \right\} + \frac{1}{2} \left\| \psi(\cdot) \right\|^2_{\mathrm{RKHS}}$$

- Use minimisation of empirical loss

$$\sum_{i=1}^{n} \left\{ \frac{1}{2} \left\| \nabla \ln q_\theta(x_i) \right\|^2 + \nabla^2 \ln q_\theta(x_i) \right\}$$

independent of Z_θ !

- Nonparametric extension (Sriperumbudur, Fukumizu, Kumar, Gretton and Hyvärinen, 2014). Set $\psi(x) \doteq \ln \tilde{q}(x)$

$$\sum_{i=1}^{n} \left\{ \frac{1}{2} \left\| \nabla \psi(x_i) \right\|^2 + \nabla^2 \psi(x_i) \right\} + \frac{1}{2} \left\| \psi(\cdot) \right\|^2_{\mathrm{RKHS}}$$

- yields estimate of $\nabla \ln p(x)$!

Some applications of score matching

- Learning structure of graphical models
- Finding modes of probability densities
- Gradient free Hamiltonian Monte Carlo
- Sequential importance sampling
- ...

Outline

- Problem of learning drift functions for stochastic differential equations
- Nonparametric estimates using a generalisation of score matching
- Applications to Langevin models
- Relation to Maximum likelihood and Bayes
- Kullback–Leibler divergence, control and a normalizer
- Future work and open problems

Stochastic differential equations

- Dynamics defined by SDEs for $Z \in R^d$.

$$dZ_t = \underbrace{g(Z_t)}_{\text{Drift}} dt + \underbrace{\sigma(Z_t)}_{\text{Diffusion}} \times \underbrace{dW_t}_{\text{Wiener process}}$$

- Limit of discrete time process

$$Z_{t+\Delta} - Z_t = g(Z_t)\Delta + \sigma(Z_t)\sqrt{\Delta}\,\epsilon_t \ .$$

with ϵ_t i.i.d. Gaussian for $\Delta \to 0$.

Stochastic differential equations

- Dynamics defined by SDEs for $Z \in R^d$.

$$dZ_t = \underbrace{g(Z_t)}_{\text{Drift}} dt + \underbrace{\sigma(Z_t)}_{\text{Diffusion}} \times \underbrace{dW_t}_{\text{Wiener process}}$$

- Limit of discrete time process

$$Z_{t+\Delta} - Z_t = g(Z_t)\Delta + \sigma(Z_t)\sqrt{\Delta}\, \epsilon_t \ .$$

  with ϵ_t i.i.d. Gaussian for $\Delta \to 0$.

- Learn the function $g(\cdot)$ from a set of (noise free) observations $z(t_1), z(t_2), \ldots, z(t_n)$.

Nonparametric (Gaussian process) approach

Use a Gaussian Process prior $g(\cdot) \sim \mathcal{GP}(0, K(z, z'))$ over drift functions. (Papaspilioupoulis, Pokern, Roberts and Stuart 2012).

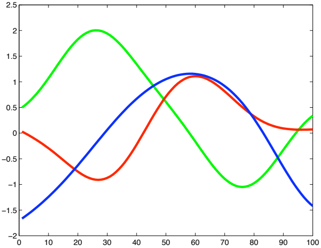

Likelihood for densely observed path

- In Euler discretization the SDE looks like $Z_{t+\Delta t} - Z_t = g(Z_t)\Delta + \sqrt{\Delta}\,\epsilon_t$, for $\Delta \to 0$.

Likelihood for densely observed path

- In Euler discretization the SDE looks like $Z_{t+\Delta t} - Z_t = g(Z_t)\Delta + \sqrt{\Delta}\,\epsilon_t$, for $\Delta \to 0$.
- Hence the likelihood for the drift is

$$p(Z_{0:T}|g) \propto \exp\left[-\frac{1}{2\Delta}\sum_t ||Z_{t+\Delta} - Z_t||^2\right] \times$$
$$\exp\left[-\frac{1}{2}\sum_t ||g(Z_t)||^2\,\Delta + \sum_t g(Z_t)\cdot(Z_{t+\Delta} - Z_t)\right].$$

allows for simple GP based estimation of the function $g(\cdot)$.

Likelihood for densely observed path

- In Euler discretization the SDE looks like
  $Z_{t+\Delta t} - Z_t = g(Z_t)\Delta + \sqrt{\Delta}\,\epsilon_t$, for $\Delta \to 0$.
- Hence the likelihood for the drift is

$$p(Z_{0:T}|g) \propto \exp\left[-\frac{1}{2\Delta}\sum_t ||Z_{t+\Delta} - Z_t||^2\right] \times$$

$$\exp\left[-\frac{1}{2}\sum_t ||g(Z_t)||^2\,\Delta + \sum_t g(Z_t)\cdot(Z_{t+\Delta} - Z_t)\right].$$

  allows for simple GP based estimation of the function $g(\cdot)$.
- This essentially leads to the estimate
  $g(z) \approx E\left[\frac{Z_{t+\Delta}-Z_t}{\Delta}|Z_t = z\right]$. Works well for $\Delta \to 0$.

For not so small Δ it does not work well ! Data from $dz = (z - z^3)dt + dW$.

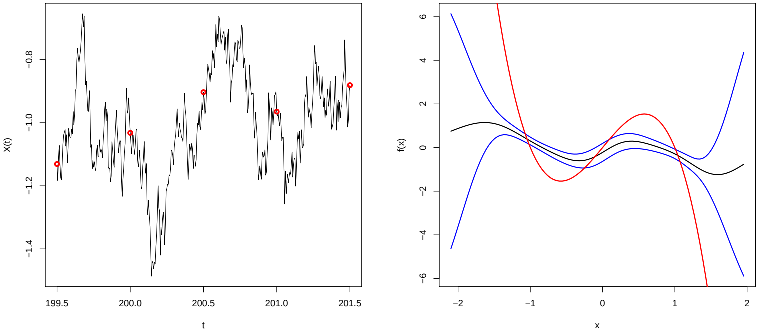

Estimation from sparse observations in time not trivial ! Approximation using imputation of hidden process possible (Ruttor, Batz and Opper, 2013).

Drift estimation from empirical density only

- Given stationary density $p(z)$ of the process. Determine the drift g ? Assume that $\sigma(\cdot)$ is known and $g(z) = r(z) + f(z)$, with $r(z)$ known.

Drift estimation from empirical density only

- Given stationary density $p(z)$ of the process. Determine the drift g ?
  Assume that $\sigma(\cdot)$ is known and $g(z) = r(z) + f(z)$, with $r(z)$ known.
- Partial answer:
  'minimal' solution minimising a quadratic functional

$$\frac{1}{2}\int p(z)\; f(z)\cdot A^{-1}(z) f(z)\; dz$$

Drift estimation from empirical density only

- Given stationary density $p(z)$ of the process. Determine the drift g ? Assume that $\sigma(\cdot)$ is known and $g(z) = r(z) + f(z)$, with $r(z)$ known.
- Partial answer:
  'minimal' solution minimising a quadratic functional

$$\frac{1}{2}\int p(z)\, f(z) \cdot A^{-1}(z) f(z)\, dz$$

- Lagrange–functional

$$\frac{1}{2}\int f(z) \cdot A^{-1}(z) f(z)\, dz - \int \psi(z) \left\{ \mathcal{L}p(z) - \nabla \cdot (f(z)p(z)) \right\}\, dz$$

  with Fokker–Planck operator corresponding to $r(z)$

$$\mathcal{L}p(z) = -\nabla \cdot (r(z)p(z)) + \frac{1}{2}\mathrm{tr}\left[\nabla\nabla^\top (D(z)p(z))\right]$$

  with $D(z) \doteq \sigma(z)\sigma(z)^\top$.

Variation yields $f(z) = A(z)\nabla\psi(z)$. Inserting back into Lagrangean yields dual functional

$$\varepsilon[\psi] = \int \left\{ \frac{1}{2}\nabla\psi(z) \cdot A(z)\ \nabla\psi(z) + \mathcal{L}^*\psi(z) \right\} p(z)dz$$

with $\mathcal{L}^*$ adjoint operator which fulfils $\int \psi(z)\mathcal{L}p(z)dz = \int p(z)\mathcal{L}^*\psi(z)dz$ and is given by

$$\mathcal{L}^*\psi(z) = r(z) \cdot \nabla\psi(z) + \frac{1}{2}\mathrm{tr}\left[D(z)\nabla\nabla^\top\psi(z) \right]$$

For 'thermal equilibrium' $A = I$ and $D = 2I$ and $r = 0$ this corresponds to score matching ! Stationary density $p(z) \propto e^{2\psi(z)}$ and $f(z) = \nabla\psi(x)$

Regularized empirical loss

- Given ergodic sample $\{z_i\}_{i=1}^n$ replace

$$p(z) \to \hat{p}(z) = \frac{1}{n} \sum_{i=1}^{n} \delta(z - z_i)$$

- Results in empirical functional

$$C \sum_{i=1}^{n} \left\{ \frac{1}{2} \nabla\psi(z_i) \cdot A(z_i)\ \nabla\psi(z_i) + \mathcal{L}^* \psi(z_i) \right\}$$

Regularized empirical loss

- Given ergodic sample $\{z_i\}_{i=1}^n$ replace

$$p(z) \to \hat{p}(z) = \frac{1}{n}\sum_{i=1}^{n} \delta(z - z_i)$$

- Results in empirical functional

$$C \sum_{i=1}^{n} \left\{ \frac{1}{2} \nabla\psi(z_i) \cdot A(z_i)\, \nabla\psi(z_i) + \mathcal{L}^* \psi(z_i) \right\} + \frac{1}{2} \int\int \psi(z) K^{-1}(z, z') \psi(z')\, dz dz'$$

regularised by kernel K.

- Variation wrt ψ yields

$$\psi(z) + C \sum_{j=1}^{n} \mathcal{L}^*_{z'}[\psi]\ K(z, z')_{z'=z_j} = 0$$

- Variation wrt ψ yields

$$\psi(z) + C \sum_{j=1}^{n} \mathcal{L}^*_{z'}[\psi] \; K(z, z')_{z'=z_j} = 0$$

- regularised version of equation (valid for any function $h(\cdot)$)

$$\int p(z')\mathcal{L}^*_{z'}[\psi]h(z')dz' = \int h(z')\mathcal{L}_{z'}[\psi]p(z')dz' = 0$$

applied to $p \to \hat{p}$ and $h_z(z') = K(z, z')$.

- Variation wrt ψ yields

$$\psi(z) + C \sum_{j=1}^{n} \mathcal{L}^*_{z'}[\psi]\; K(z, z')_{z'=z_j} = 0$$

- regularised version of equation (valid for any function $h(\cdot)$)

$$\int p(z')\mathcal{L}^*_{z'}[\psi]h(z')dz' = \int h(z')\mathcal{L}_{z'}[\psi]p(z')dz' = 0$$

applied to $p \to \hat{p}$ and $h_z(z') = K(z, z')$.

- If $\nabla\psi(z)$ known at all sample points $z = z_i$, get $\psi(z)$ for all z.
- Take gradient

$$\nabla\psi(z_i) + C \sum_{j=1}^{n} \mathcal{L}^*_{z'}[\psi]\nabla_z K(z, z')_{z=z_i, z'=z_j} = 0$$

Langevin dynamics

- Classical mechanics in terms of (generalized) coordinates and velocities $X, V \in R^d$

$$dX_t = V_t dt, \qquad dV_t = g_v(X_t, V_t)dt + \sigma_v(X_t, V_t)dW_t.$$

Langevin dynamics

- Classical mechanics in terms of (generalized) coordinates and velocities $X, V \in R^d$

$$dX_t = V_t dt, \qquad dV_t = g_v(X_t, V_t)dt + \sigma_v(X_t, V_t)dW_t.$$

- If drift is of the form $g_v(x, v) = r_v(x, y) + \nabla_v \psi(x, v)$ use functional for estimation

$$\varepsilon[\psi] = \int p(x, v) \left\{ \mathcal{L}^* \psi(x, v) + \frac{1}{2}(\nabla_v \psi(x, v))^2 \right\} dx\, dv$$

where

$$\mathcal{L}^* \psi(x, v) = \left(v \cdot \nabla_x + r_v(x, v) \cdot \nabla_v + \frac{1}{2}\mathrm{tr}(D_v(x, v)\nabla_v^\top \nabla_v) \right) \psi(x, v) \, .$$

- Condition $f_v(x, v) = \nabla_v \phi(x, v)$ restricts the velocity dependency!

- Condition $f_v(x, v) = \nabla_v \phi(x, v)$ restricts the velocity dependency!
- Possible choice

$$f_v(x, v) = f(x) - \Lambda v = \nabla_v \left\{ v \cdot f(x) - \frac{1}{2} v \cdot \Lambda v \right\} ,$$

- Condition $f_v(x, v) = \nabla_v \phi(x, v)$ restricts the velocity dependency!
- Possible choice

$$f_v(x, v) = f(x) - \Lambda v = \nabla_v \left\{ v \cdot f(x) - \frac{1}{2} v \cdot \Lambda v \right\} \ ,$$

- Note: If $r_v = -\Lambda v$ then $\mathcal{L}^* \psi(x, v) = \mathcal{L}^*(v \cdot f(x))$ is independent of the diffusion term $D_v(x, v)$:
  $\rightarrow$ Estimate $f(x)$ without knowing the diffusion.

Example: Double well

$$f(x, v) = 4(x - x^3) - \lambda v$$

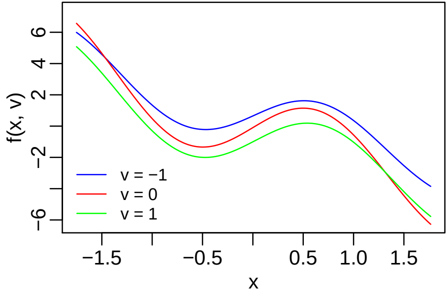

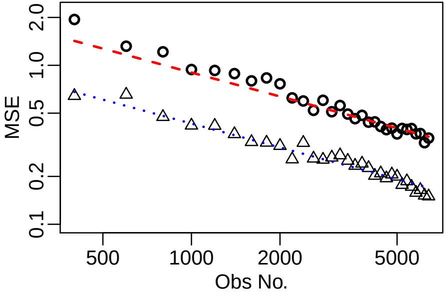

Example: Nonconservative force

$$f^{(1)}(x) = x^{(1)}(1 - (x^{(1)})^2 - (x^{(2)})^2) - x^{(2)}$$
$$f^{(2)}(x) = x^{(2)}(1 - (x^{(1)})^2 - (x^{(2)})^2) - x^{(1)}$$

Polynomial kernel with $p = 4$ and $n = 2000$

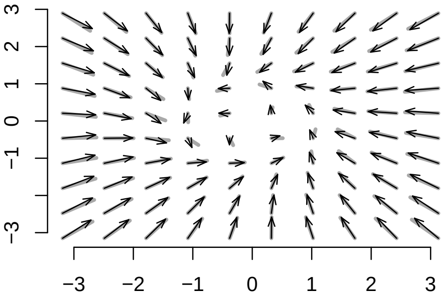

Cart and Pole model

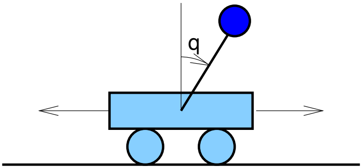

$f(x) = a \sin x$ and $r(v) = -\lambda v$ and diffusion $D_v = (\sigma \cos(x))^2$.

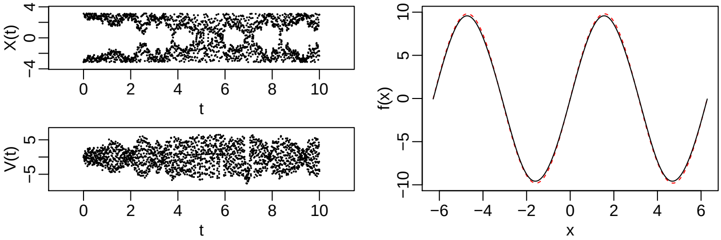

Kernel density estimators as alternative ?

Explicit solutions to drift for Langevin models

$$f^{(i)}(x) = \sum_{j=1}^{d} \frac{\partial E[v^{(i)} v^{(j)} | x]}{\partial x^{(j)}} + \sum_{j=1}^{d} E[v^{(i)} v^{(j)} | x] \frac{\partial \ln p(x)}{\partial x^{(j)}} - E[r^{(i)} | x] \ ,$$

Kernel density estimators as alternative ?

Explicit solutions to drift for Langevin models

$$f^{(i)}(x) = \sum_{j=1}^{d} \frac{\partial E[v^{(i)} v^{(j)}|x]}{\partial x^{(j)}} + \sum_{j=1}^{d} E[v^{(i)} v^{(j)}|x] \frac{\partial \ln p(x)}{\partial x^{(j)}} - E[r^{(i)}|x] \; ,$$

Simplifies only for 'thermal equilibrium' $f(x) = \nabla\phi$ and $D_v \propto \Sigma$, $r = -\Lambda v$ where Λ and Σ are diagonal with $\frac{2\lambda_i}{\sigma_i^2} = \beta$. One has then $E[v^{(i)} v^{(j)}|x] = \frac{1}{2\beta}\delta_{ij}$ and $E[r^{(i)}|x] = 0$.

Extension: Other evolution equations

- Replace white noise $\sigma_v(X, V)dW \to U(t)dt$ where $U(t)$ Markovian.
- Include noise in state variable $Z = (X, V, U)$
- Example: $U(t) = \pm 1$ random telegraph process and $x, v, u \in R$ with drift $g_v(x, v) = -\lambda v + f(x)$

Extension: Other evolution equations

- Replace white noise $\sigma_v(X, V)dW \to U(t)dt$ where $U(t)$ Markovian.
- Include noise in state variable $Z = (X, V, U)$
- Example: $U(t) = \pm 1$ random telegraph process and $x, v, u \in R$ with drift $g_v(x, v) = -\lambda v + f(x)$
- Fokker Planck $\to$ Master equation

$$
\begin{aligned}
0 = -\partial_x(vp(x, v, u) - \partial_v \left[(f(x) - \lambda v + u)p(x, v, u)\right] \\
+\gamma \left(p(x, v, -u) - p(x, v, u)\right)
\end{aligned}
$$

Extension: Other evolution equations

- Replace white noise $\sigma_v(X, V)dW \to U(t)dt$ where $U(t)$ Markovian.
- Include noise in state variable $Z = (X, V, U)$
- Example: $U(t) = \pm 1$ random telegraph process and $x, v, u \in R$ with drift $g_v(x, v) = -\lambda v + f(x)$
- Fokker Planck $\to$ Master equation

$$0 = -\partial_x (vp(x, v, u) - \partial_v \left[(f(x) - \lambda v + u) p(x, v, u) \right] \\ + \gamma \left(p(x, v, -u) - p(x, v, u) \right)$$

- Adjoint operator

$$\mathcal{L}^* \psi = \{ v \partial_x + (u - \lambda v) \partial_v \} \psi + \gamma \left(\psi(x, v, -1) - \psi(x, v, 1) \right)$$

Extension: Other evolution equations

- Replace white noise $\sigma_v(X, V)dW \to U(t)dt$ where $U(t)$ Markovian.
- Include noise in state variable $Z = (X, V, U)$
- Example: $U(t) = \pm 1$ random telegraph process and $x, v, u \in R$ with drift $g_v(x, v) = -\lambda v + f(x)$
- Fokker Planck $\to$ Master equation

$$0 = -\partial_x(vp(x, v, u) - \partial_v \left[(f(x) - \lambda v + u)p(x, v, u)\right] \\ +\gamma\left(p(x, v, -u) - p(x, v, u)\right)$$

- Adjoint operator

$$\mathcal{L}^*\psi = \left\{v\partial_x + (u - \lambda v)\partial_v\right\}\psi + \gamma\left(\psi(x, v, -1) - \psi(x, v, 1)\right)$$

- Parametrisation $\psi(x, v) = vf(x)$ leads to functional

$$\varepsilon[f] = \frac{1}{2}\sum_{u=\pm 1}\int p(x, v, u)\left\{f^2(x) + 2f'(x)v^2 + 2f(x)(u - \lambda v)\right\} dxdv$$

The case $A = D$: Likelihood for densely observed path

$$-\ln p(Z_{0:T}|g) = \frac{1}{2}\sum_t \left\{||g(Z_t)||^2 \Delta t - 2\langle g(Z_t), (Z_{t+\Delta t} - Z_t)\rangle\right\} + \text{const}$$

with $\langle u, v\rangle \doteq u^\top D^{-1} v$.

The case $A = D$: Likelihood for densely observed path

$$-\ln p(Z_{0:T}|g) = \frac{1}{2}\sum_t \left\{||g(Z_t)||^2 \Delta t - 2\langle g(Z_t), (Z_{t+\Delta t} - Z_t)\rangle\right\} + \text{const}$$

with $\langle u, v\rangle \doteq u^\top D^{-1} v$. Assume $g = r + D\nabla\psi$ take $\Delta t \to 0$ and apply Ito formula

$$= \text{const} + \frac{1}{2}\int_0^T \{\nabla\psi \cdot D\ \nabla\psi\, dt + 2r \cdot \nabla\psi\, dt - 2\nabla\psi \cdot dZ_t\}$$

The case $A = D$: Likelihood for densely observed path

$$-\ln p(Z_{0:T}|g) = \frac{1}{2}\sum_t \left\{||g(Z_t)||^2 \Delta t - 2\langle g(Z_t), (Z_{t+\Delta t} - Z_t)\rangle\right\} + \text{const}$$

with $\langle u, v\rangle \doteq u^\top D^{-1} v$. Assume $g = r + D\nabla\psi$ take $\Delta t \to 0$ and apply Ito formula

$$= \text{const} + \frac{1}{2}\int_0^T \{\nabla\psi \cdot D\ \nabla\psi\, dt + 2r \cdot \nabla\psi\, dt - 2\nabla\psi \cdot dZ_t\}$$

$$= \quad \frac{1}{2}\int_0^T \left\{\nabla\psi \cdot D\ \nabla\psi + 2r \cdot \nabla\psi + \mathrm{tr}(D\nabla\nabla^\top\psi)\right\} dt - \psi(Z_T) + \psi(Z_0)$$

The case $A = D$: Likelihood for densely observed path

$$-\ln p(Z_{0:T}|g) = \frac{1}{2}\sum_t \left\{||g(Z_t)||^2 \Delta t - 2\langle g(Z_t), (Z_{t+\Delta t} - Z_t)\rangle\right\} + \text{const}$$

with $\langle u, v\rangle \doteq u^\top D^{-1} v$. Assume $g = r + D\nabla\psi$ take $\Delta t \to 0$ and apply Ito formula

$$= \text{const} + \frac{1}{2}\int_0^T \{\nabla\psi \cdot D\, \nabla\psi\, dt + 2r \cdot \nabla\psi\, dt - 2\nabla\psi \cdot dZ_t\}$$

$$= \frac{1}{2}\int_0^T \left\{\nabla\psi \cdot D\, \nabla\psi + 2r \cdot \nabla\psi + \mathrm{tr}(D\nabla\nabla^\top\psi)\right\} dt - \psi(Z_T) + \psi(Z_0)$$

$$\simeq \frac{T}{2}\int \left\{\nabla\psi \cdot D\, \nabla\psi + 2r \cdot \nabla\psi + \mathrm{tr}(D\nabla\nabla^\top\psi)\right\} p(z)dz$$

Back to the optimisation problem

We considered minimisation of

$$\frac{1}{2}\int p(z)\ f(z)\cdot A^{-1}(z)f(z)\ dz$$

under the linear constraint that

$$-\nabla\cdot(r(z)p(z))+\frac{1}{2}\mathrm{tr}\left[\nabla\nabla^{\top}(D(z)p(z))\right]=0$$

The case $A = D$: Kullback Leibler divergence

- Kullback-Leibler (KL) divergence between the path probabilities for diffusion processes with drifts $g(z) = r(z) + f(z)$ and $r(z)$:

$$D(p(Z_{0:T}|r+f)||p(Z_{0:T}|r)) = \frac{1}{2}\int_0^T dt \int p_t(z) f(z) \cdot D^{-1}(z) f(z) dz$$

assuming equal diffusions $D(z)$.

The case $A = D$: Kullback Leibler divergence

- Kullback-Leibler (KL) divergence between the path probabilities for diffusion processes with drifts $g(z) = r(z) + f(z)$ and $r(z)$:

$$D(p(Z_{0:T}|r+f)||p(Z_{0:T}|r)) = \frac{1}{2}\int_0^T dt \int p_t(z)f(z)\cdot D^{-1}(z)f(z)dz$$

  assuming equal diffusions $D(z)$.

- For $T \to \infty$ the relative entropy rate is

$$\lim_{T\to\infty}\frac{1}{T}D(p(Z_{r+f}0:T)||p_r(Z_{0:T}|r) = \frac{1}{2}\int p(z)f(z)\cdot D^{-1}(z)f(z)dz.$$

  Hence, the choice $A = D$ may be understood as a generalized as **minimum relative entropy** solution where the stationary density is given as a constraint.

KL control problem and normalizer

- Find extra drift (control) $f(\cdot)$ such that the cost rate

$$C[f] \doteq \lim_{T\to\infty} \frac{1}{T} D(p_{r+f}(Z_{0:T})||p_r(Z_{0:T}) + \int p_{r+f}(z)U(z)dz$$

is minimized ! $U(z)$ are state costs.

KL control problem and normalizer

- Find extra drift (control) $f(\cdot)$ such that the cost rate

$$C[f] \doteq \lim_{T\to\infty} \frac{1}{T} D(p_{r+f}(Z_{0:T})||p_r(Z_{0:T}) + \int p_{r+f}(z)U(z)dz$$

is minimized ! $U(z)$ are state costs.

- Note that for the finite T problem the controlled path probabilities satisfies

$$p_{r+f}(Z_{0:T} = \frac{1}{\zeta_T} p_r(Z_{0:T})e^{-\int_0^T U(Z_t)dt}$$

where the normaliser is given by

$$\zeta_T = E_r\left[e^{-\int_0^T U(Z_t)dt}\right] \qquad (1)$$

KL control problem contd

If the stationary controlled density p is known (or we can sample from it) we can compute

- the log–normalizer in the limit via

$$- \lim_{T \to \infty} \frac{1}{T} \ln \zeta_T = \int p(z) U(z) dz -$$

$$\min_{\psi} \int \left\{ \frac{1}{2} \nabla \psi(z) \cdot A(z) \, \nabla \psi(z) + \mathcal{L}^* \psi(z) \right\} p(z) dz$$

KL control problem contd

If the stationary controlled density p is known (or we can sample from it) we can compute

- the log–normalizer in the limit via

$$-\lim_{T\to\infty}\frac{1}{T}\ln\zeta_T = \int p(z)U(z)dz -$$
$$\min_{\psi}\int\left\{\frac{1}{2}\nabla\psi(z)\cdot A(z)\ \nabla\psi(z) + \mathcal{L}^*\psi(z)\right\}p(z)dz$$

- and the control via

$$f(z) = D(z)\nabla\ln\psi(z)$$

Could this be helpful to solve the control problem ?

- Sampling from the 'smoothing density' $p(z)$ is not easy.
- Sampling from the 'filtering density' $p_t^{filt}(z) \doteq p(z|U(Z_\tau), \tau \leq t)$ is possible using particle filters.
- In certain cases estimates of $\nabla \ln p_t^{filt}(z)$ can be related to $\nabla \ln p(z)$

Analytically solvable example: Uncontrolled drift $r(z) = 4z(1 - z^2)$ and path costs $U(z) = \beta r(z)$.

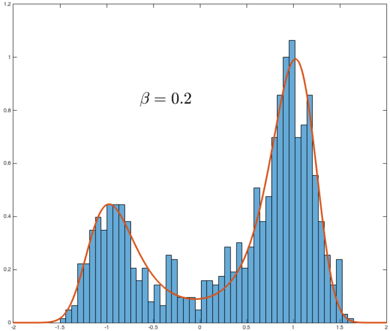

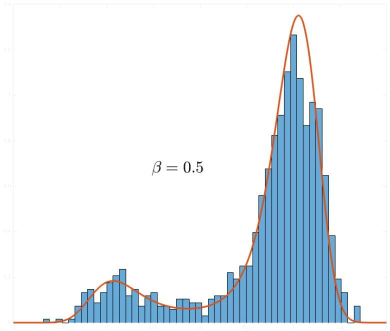

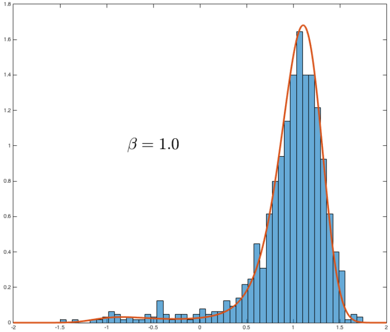
$\beta = 1.0$
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

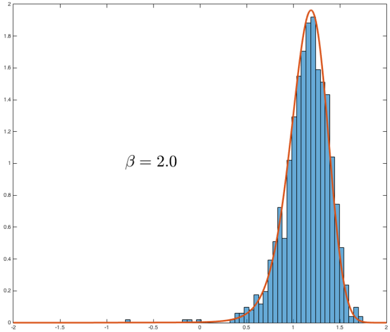
$\beta = 2.0$
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

Future work and open problems

- Generalisation to dynamics
- Other Markov processes
- Relation to Bayes and hyper parameter selection
- Convergence rates for non i.i.d. data
- Inclusion of noise in data ?
- Unobserved variables ?
- Inclusion of information from time order of data ?