

Least-action filtering

L. C. G. Rogers
University of Cambridge
First version: September 2008

April 23, 2012

Abstract

This paper presents an approach to estimating a hidden process in a continuous-time setting, where the hidden process is a diffusion. The approach is simply to minimize the negative log-likelihood of the hidden path, where the likelihood is expressed relative to Wiener measure. This negative log-likelihood is the action integral of the path, which we minimize by calculus of variations. We then perform an asymptotic maximum-likelihood analysis to understand better how the actual path is distributed around the least-action path; it turns out that the actual path can be expressed (approximately) as the sum of the least-action path and a zero-mean Gaussian process which can be specified quite explicitly. Numerical solution of the ODEs which arise from the calculus of variations is often feasible, but is complicated by the shooting nature of the problem, and the possibility that we have found a local but not global minimum. We analyze the situations when this happens, and provide effective numerical methods for studying this. We also show how the methodology works in a situation where the hidden positive diffusion acts as the random intensity of a point process which is observed; here too it is possible to estimate the hidden process.

1 Introduction.

The basic problem tackled in this paper is to try to estimate the hidden part $(X_t)_{0 \leq t \leq T}$ of a vector¹ diffusion process $Z_t \equiv [X_t; Y_t]$ given the observations $(Y_t)_{0 \leq t \leq T}$. Of course, this is an idealized question, because in practice we would only observe the signal at discrete instants of time, which for simplicity we will assume are equally spaced with a spacing of $h > 0$. It would then in principle be possible to write down the likelihood² of $(Z_{nh})_{0 \leq n \leq N}$, and maximize this over the unobserved variables $(X_{nh})_{0 \leq n \leq N}$, with the Y -values known and fixed.

¹We use the Scilab/Matlab notation, where $z = [x; y]$ denotes the column vector z formed by stacking the column vector x above the column vector y .

²We suppose that $T = Nh$.

There are two quite substantial obstacles on this path. The first is that (except in a few special examples) there is no closed-form expression for the transition density of a general diffusion, so evaluating the likelihood is already problematic. The second is that we are faced with a maximization over a space of dimension proportional to N , a number which will be large if h is small, as we envisage. The approach of this paper avoids these difficulties completely by passing to a limiting form of the problem in which the (negative of the) log-likelihood expression to be maximized converges to an integral, the action integral. This can be minimized by calculus of variations, so that instead of having to do some numerical nonlinear optimization, we find ourselves having to solve a second-order non-linear ordinary differential equation (ODE), which is a far easier task. The only issue to be dealt with is that the ODE is of shooting type, with boundary conditions at $t = 0$ and $t = T$, so that some iterative solution scheme is needed: the analysis is explained in Section 3.

The action integral is only the negative log-likelihood of the path in some formal sense, since it involves derivatives of the paths of the diffusion, and diffusion paths are not differentiable. Nevertheless, it is a useful heuristic³ which leads to a rigorous approximation result for the limiting form of the maximum-likelihood path as $h \downarrow 0$: see Theorem 1.

This gives us a way to identify effectively the ‘most likely’ path x^* of X given the path of Y , but we would also like to have some idea of how the random path of X is distributed around x^* . This involves a study of the log-likelihood in a neighbourhood of x^* which we find is (to leading order) quadratic in the perturbation $\xi = x - x^*$. Thus the perturbation ξ is approximately a Gaussian process, whose mean is identically zero, and whose law can be precisely characterized by expressing ξ as the solution of a linear stochastic differential equation: see Section 4.

Inference on a hidden diffusion intensity for a point process is discussed in Section 5, and the relationship with SMC methodologies is discussed briefly in Section 6. Section 7 concludes.

The idea of studying the action of a diffusion path is quite ancient already, and even as a tool for estimation there is a literature: see, for example, [4], [1], [2]. Markussen [4] arrives at essentially the same identification of the maximum-likelihood path as we do; the expression derived here for the Gaussian law of the perturbation is considerably more explicit. The approach of Archambeau et al [1] has superficial similarities, but is different in major respects; in particular, they use a minimum relative entropy criterion to identify a ‘best’ Gaussian approximation to the hidden process, whose structure appears to require the calculation of expectations of non-linear functionals of the path. In contrast, the approach followed here requires only the solution of ordinary differential equations⁴ This is particularly advantageous in higher dimensions, as the numerical solution of ordinary differential equations still works quite effectively in moderately large dimension, whereas methods such as particle filtering

³This heuristic leads in a few lines to (non-rigorous) derivations of the Cameron-Martin-Girsanov change of measure, and to large-deviation rate functions for various diffusion asymptotics.

⁴Archambeau et al work with a more restricted diffusion dynamic (which they assert can be extended), and a somewhat different structure for the relation between observations and the hidden process.

struggle in dimension much above about seven, as we shall discuss in Section 6.

2 The setting.

Suppose that we have some stochastic differential equation (SDE) in \mathbb{R}^d

$$dZ_t = \sigma(t, Z_t) dW_t + \mu(t, Z_t) dt, \quad (1)$$

where Z is partitioned $Z = [X; Y]$ and where X is n -dimensional, Y is s -dimensional, $n+s = d$. We shall suppose also that σ , σ^{-1} and μ are in C_b^1 . The problem we address is to estimate the hidden part $(X_t)_{0 \leq t \leq T}$ of Z given observations of $(Y_t)_{0 \leq t \leq T}$; the possibility that Y is non-informative about X is included in the analysis, so what follows applies to the distribution of an unobserved diffusion also.

Closely related to (1) is the first-order Euler-Maruyama difference scheme

$$dz_t^{(n)} = \sigma(t_n, z_{t_n}^{(n)}) dW_t + \mu(t_n, z_{t_n}^{(n)}) dt \quad (2)$$

where $t_n \equiv 2^{-n}[2^n t]$. Despite appearances, this can be viewed as a discrete scheme; the increments of z have conditionally Gaussian distributions. It is known (see, for example, Mao [3]) that for some constant C

$$E \left[\sup_{0 \leq t \leq T} |Z_t - z_t^{(n)}|^2 \right] \leq C 2^{-n}, \quad (3)$$

so by the first Borel-Cantelli Lemma we conclude that there is almost-sure uniform convergence of the processes $z^{(n)}$ to the solution Z to the SDE.

In the discrete-time approximation (2) to the SDE, the conditional density of $x(j2^{-n})_{0 \leq j \leq 2^n T}$ given $y(j2^{-n})_{0 \leq j \leq 2^n T}$ can be written down immediately. The log-likelihood is

$$\begin{aligned} \lambda_n(x|y) &= -\frac{1}{2} \sum_{j=0}^{N-1} \frac{1}{h} \left| \sigma(jh, z_{jh})^{-1} (z_{jh+h} - z_{jh} - h\mu(jh, z_{jh})) \right|^2 - \varphi(x_0) \\ &= -\frac{1}{2} \sum_{j=0}^{N-1} h \left| \sigma(jh, z_{jh})^{-1} \left(\frac{z_{jh+h} - z_{jh}}{h} - \mu(jh, z_{jh}) \right) \right|^2 - \varphi(x_0), \end{aligned} \quad (4)$$

where the prior density of x_0 is $\exp(-\varphi(x))$, and $h = 2^{-n}$. Inspecting (4), we see a difference quotient of z which it is tempting to replace with a derivative as $n \rightarrow \infty$, leading to the *formal* limit

$$\Lambda(x|y) = -\frac{1}{2} \int_0^T \left| \sigma(s, z_s)^{-1} (\dot{z}_s - \mu(s, z_s)) \right|^2 ds - \varphi(x_0) \quad (5)$$

$$\equiv - \int_0^T \psi(s, x_s, p_s) ds - \varphi(x_0) \quad (6)$$

where the Riemann sum becomes an integral, and we write $p_s \equiv \dot{x}_s$. This is a perfectly sensible functional of a C^1 path z , even though for a diffusion process the path will not be differentiable.

Our viewpoint is that the problem which concerns us is to estimate $x(j2^{-\nu})_{0 \leq j \leq 2^\nu T}$ in the *discretely-sampled* model (2), for some *fixed* (quite large) integer ν , and that the SDE version of the problem is to be used to help us in this. In practice, we will only ever have the observation data $y(j2^{-\nu})_{0 \leq j \leq 2^\nu T}$ in discretely-sampled form (how would it be stored if not?!) so this is the realistic question. When we minimize the log-likelihood (4) over x , we only need the values of y at the multiples of $h = 2^{-\nu}$; we therefore lose no generality in supposing that y has been interpolated (by cubic splines, say) to be C^2 in all of $[0, T]$. This does not of course change (4), but it does mean that the functional Λ is meaningfully defined for any C^1 path $(x_t)_{0 \leq t \leq T}$. Given this, we have the following result, where to emphasize again, we are considering what happens as $n \geq \nu$ tends to infinity, with ν fixed.

Theorem 1 *Suppose that σ, σ^{-1} and μ are C_b^1 , and that $(y_t)_{0 \leq t \leq T}$ is also C^1 . Suppose further that*

$$\lim_{|x| \rightarrow \infty} \varphi(x) = \infty. \quad (7)$$

Then

$$\lim_{n \rightarrow \infty} \inf_x \{ -\lambda_n(x|y) \} = \inf_x \{ -\Lambda(x|y) \}. \quad (8)$$

Assuming that $\Lambda(\cdot|y)$ has a unique maximizer x^ , and that x_n is a maximizer of $\lambda_n(\cdot|y)$, then $x_n \rightarrow x^*$ uniformly.*

PROOF. See Appendix A.

The usefulness of Theorem 1 is that it shows that a minimizer of $-\lambda_n(\cdot|y)$ is very close to the (assumed unique) minimizer of $-\Lambda(\cdot|y)$; but this minimizer can be found by calculus of variations without resort to some high-dimensional non-linear optimizer.

Notice that although we make a discretization error when we replace the SDE (1) with the Euler-Maruyama scheme (2), in most examples of practical interest the magnitude of this discretization error will be insignificant compared to the observation error that we are attempting to see past.

3 Finding the least-action path.

Theorem 1 shows that asymptotically as the step size $h = 2^{-n}$ tends to zero, the maximum-likelihood estimate of the unobserved path converges to the maximizer of the continuous-time analogue Λ . Resorting now to calculus of variation techniques leads us to the following central result.

Theorem 2 *The path (x, p) which minimizes the action functional*

$$-\Lambda(x|y) = \int_0^T \psi(s, x_s, p_s) ds + \varphi(x_0) \quad (9)$$

satisfies the following ODE with boundary conditions:

$$0 = D_{p_j} \psi(0, x_0, p_0) - D_{x_j} \varphi(x_0) \quad (10)$$

$$0 = D_{x_j} \psi - D_t D_{p_j} \psi - \dot{x}_k D_{p_j} D_{x_k} \psi - \dot{p}_k D_{p_j} D_{p_k} \psi \quad (11)$$

$$0 = D_{p_j} \psi(T, x_T, p_T) \quad (12)$$

PROOF. The negative log-likelihood to be minimized is

$$\varphi(x_0) + \int_0^T \psi(t, x_t, \dot{x}_t) dt \equiv \varphi(x_0) + \int_0^T \psi(t, x_t, p_t) dt, \quad (13)$$

where we write $p \equiv \dot{x}$. This we attack by calculus of variations; if we have found the optimal x , then any perturbation to $x + \xi$ must to leading order make zero change to the objective. Writing down the first-order change and integrating by parts gives

$$\begin{aligned} 0 &= \xi_0 D\varphi(x_0) + \int_0^T \{ \xi_t \cdot D_x \psi + \dot{\xi} \cdot D_p \psi \} dt \\ &= \xi_0 D\varphi(x_0) + [\xi_t^j D_{p_j} \psi]_0^T + \int_0^T \xi_t^j \{ D_{x_j} \psi - D_t D_{p_j} \psi - \dot{x}_k D_{p_j} D_{x_k} \psi - \dot{p}_k D_{p_j} D_{p_k} \psi \} dt. \end{aligned} \quad (14)$$

Since ξ is arbitrary, we deduce the conditions (10), (11), (12) for optimality. \square

REMARKS. (i) We have found a (generally non-linear) first-order ODE (11) for (x, p) , with initial condition (10) and terminal condition (12). This will generally have a unique solution, though the ‘shooting’ nature of the ODE is rather clumsy in practice.

(ii) We expect that this methodology will be advantageous in higher-dimensional problems; algorithms for solving ODEs in high dimension tend to degrade less rapidly than (for example) gradient optimization methods, or particle filtering.

(iii) If the time horizon T is too large, it may be that the shooting ODE is not able to identify the initial and terminal conditions with sufficient accuracy.

To illustrate the methodology in action, we present here some simple examples.

Example 1. Here we take independent one-dimensional Ornstein-Uhlenbeck processes X, y

$$\begin{aligned} dX &= \sigma_X dW - \beta_X X dt \\ dy &= \sigma_y dW' - \beta_y Y dt \end{aligned}$$

where $\sigma_X = 1.053$, $\sigma_y = 1.0127$, $\beta_X = 0.1054$ and $\beta_y = 0.0253$. We observe $Y = X + y$. Then $Z = [X; Y]$ solves

$$\begin{aligned} dZ &= \begin{pmatrix} \sigma_X & 0 \\ \sigma_X & \sigma_y \end{pmatrix} \begin{pmatrix} dW \\ dW' \end{pmatrix} + \begin{pmatrix} -\beta_X & 0 \\ -\beta_X + \beta_y & -\beta_y \end{pmatrix} \begin{pmatrix} X \\ Y \end{pmatrix} dt \\ &\equiv \sigma \begin{pmatrix} dW \\ dW' \end{pmatrix} + A \begin{pmatrix} X \\ Y \end{pmatrix} dt, \end{aligned}$$

which is a simple linear SDE. For this example,

$$\psi(t, x, p) = \frac{1}{2} \left(\begin{pmatrix} p \\ \dot{Y}_t \end{pmatrix} - A \begin{pmatrix} x \\ Y_t \end{pmatrix} \right)^T q \left(\begin{pmatrix} p \\ \dot{Y}_t \end{pmatrix} - A \begin{pmatrix} x \\ Y_t \end{pmatrix} \right).$$

A sample path of the SDE was simulated, and the results are displayed in Figure 1. The true path is dashed in green, the noisy observation is in blue, and the least-action path is in red.

Example 2. This time we have a two-dimensional linear example, with X satisfying

$$dX_t = dW_t + \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} X_t dt \equiv dW_t + A_0 X_t dt,$$

and $Y = X + y$, where y is an independent OU process

$$dy_t = \sigma dw_t - \lambda y_t dt,$$

with $\lambda = 0.005$, $\sigma = 20$. The results of the least-action analysis are displayed in Figure 2. Once again, the estimates are quite impressive, but since the underlying dynamics are a perturbation of uniform motion in a circle, this is perhaps not so surprising.

Example 3. This is another one-dimensional example, but this time non-linear. We have dynamics

$$dX_t = \sigma_X dW_t - b \sin(aX_t) dt$$

for X , where $\sigma_X = 3$, $b = 12$ and $a = 2\pi/5$. Thus the drift tries to keep the diffusion near to multiples of 5. The observation Y is as before of the form $Y = X + y$ where y is an independent OU process

$$dy_t = dW'_t - \lambda y_t dt$$

with $\lambda = 0.05$. The results of the analysis are shown in Figure 3. As can be seen, the hidden Markov process stays close to multiples of 5, occasionally moving from one value to a neighbouring value, and as the observational error gets larger, the filtering performs less well.

Example 4. The final example is a more demanding test of the methodology. Here, we take a two-dimensional non-linear dynamic for X

$$dX_t = \Sigma dW_t - b \sin(aX_t) dt$$

where

$$\Sigma^2 = \begin{pmatrix} 0.9 & 0.27 \\ 0.27 & 0.9 \end{pmatrix} \quad (15)$$

a, b as for the previous example. The observation is *univariate*, and takes the form $Y = v \cdot X + y$ where y is an independent OU process

$$dy_t = \sigma dw_t - \lambda y_t dt$$

where $v = (1, 2)$, $\lambda = 0.05$, and $\sigma = 1$. It would be indeed remarkable if the methodology was able to unscramble the two hidden components of the diffusion from observation of just one component, and the results are shown in Fig 4; sometimes the estimate is close to the true values, sometime less so. But it should be understood that this is not a limitation of the methodology, which is after all discovering the maximum-likelihood estimator of the hidden process; the relatively poor recovery of the hidden path is because we have a small signal-to-noise ratio for this example.

4 Second-order analysis.

The first problem we focused on is maximizing the log-likelihood as given for the discretized problem by (4). The unknown in this situation is the sequence $\mathcal{X} \equiv (x_{j2^{-\nu}})_{0 \leq j \leq 2^\nu T}$, and as in classical asymptotic ML theory, having identified the ML estimator $\hat{\mathcal{X}}$ of the unknown \mathcal{X} , we can enquire about the distribution of \mathcal{X} about $\hat{\mathcal{X}}$. This amounts to understanding the second derivative of the log-likelihood at the MLE. Given suitable regularity of the likelihood, a Taylor expansion shows that the distribution of \mathcal{X} around $\hat{\mathcal{X}}$ will be approximately Gaussian, with covariance equal to the second derivative of the log-likelihood. However, we have seen that there are considerable methodological advantages in passing from the discrete problem (4) to the continuous analogue (6), and these advantages apply also at the level of the second-order terms in the expansion of the log-likelihood, as we shall now see.

Theorem 3 *Suppose that x^* is a local minimizer of the functional $\Lambda(\cdot|y)$, and define the matrix-valued $n \times n$ functions of time*

$$A_t^{ij} = D_{x_i} D_{x_j} \psi, \quad B_t^{ij} = D_{x_i} D_{p_j} \psi, \quad q_t^{ij} = D_{p_i} D_{p_j} \psi, \quad (16)$$

evaluated along the path x^ . Then:*

(i) *the ODE*

$$A_t + \dot{\theta}_t = K_t^T q_t K_t = (B_t + \theta_t) q_t^{-1} (B_t^T + \theta_t), \quad \theta_T = 0 \quad (17)$$

has a unique solution;

(ii) writing $x = x^* + \xi$, we have that

$$-\Lambda(x|y) + \Lambda(x^*|y) = Q_0(\xi) + o\left(\int_0^T (|\xi_s|^2 + |\dot{\xi}_s|^2) ds\right), \quad (18)$$

where

$$Q_0(\xi) = \frac{1}{2} \xi_0 \cdot \{ D^2\varphi(x_0^*) + \theta_0 \} \xi_0 + \int_0^T \frac{1}{2} (\dot{\xi}_t + K_t \xi_t) \cdot q_t (\dot{\xi}_t + K_t \xi_t) dt, \quad (19)$$

and where

$$K_t = q_t^{-1}(B_t^T + \theta_t). \quad (20)$$

REMARKS. (i) If we write $\dot{w}_t = q_t^{1/2}(\dot{\xi}_t + K_t \xi_t)$, then the integral term in (19) takes the form $\int_0^T |\dot{w}_t|^2 dt$, which is the action integral for standard Brownian motion. Thus informally what we learn from Theorem 3 is that the perturbation ξ ‘looks like’ a continuous zero-mean Gaussian process obtained by solving the linear SDE

$$d\xi_t = -K_t \xi_t dt + q_t^{-1/2} dW_t, \quad (21)$$

and started with initial precision $D^2\varphi(x_0^*) + \theta_0$. The covariance of ξ can thus be obtained from an Itô expansion of $\xi\xi^T$:

$$d\xi\xi^T \doteq (-K_t \xi_t \xi_t^T - \xi_t \xi_t^T K_t^T + q_t^{-1}) dt, \quad (22)$$

where \doteq signifies that the two sides differ by a (local) martingale. Hence

$$\dot{V}_t = -K_t V_t - V_t K_t^T + q_t^{-1} \quad (23)$$

and this allows us to calculate the covariance at time t , again by solving an ODE.

(ii) The ODE (17) for θ is quadratic, so there is no general result which guarantees that it has a solution; the solution may explode. However, in this situation we can be sure that it will not, because $\Lambda(\cdot|y)$ is minimized at x^* , as we shall see in the proof.

(iii) It has to be admitted that the relationship between Theorem 3 and the original discretized problem is somewhat tenuous. With effort it would no doubt be possible to establish an analogue of Theorem 1 for the second-order effect, but it is harder to see what the use of such a result would be. The way we envisage using Theorem 3 would be to generate an approximate distribution for the posterior of \mathcal{X} given y , which could be checked numerically against a full posterior⁵, and which could be used to give approximate moments of the posterior law of \mathcal{X} .

⁵It would be rare that this could be found in closed form, so we would typically need to resort to particle filtering to make a numerical approximation.

PROOF OF THEOREM 3. Let us begin by assuming the first statement (i) of the theorem and showing how this leads to the more interesting statement (ii). We will return to prove (i) subsequently.

If we go back to the action functional (9) and consider expanding around the local minimizer x^* by perturbing to $x^* + \xi$ for small ξ , the second-order contribution we get is

$$\begin{aligned} Q(\xi) &\equiv \frac{1}{2} D_{x_i} D_{x_j} \varphi(x_0^*) \xi_0^i \xi_0^j + \int_0^T \left\{ \frac{1}{2} \xi^i \xi^j D_{x_i} D_{x_j} \psi + \xi^i \dot{\xi}^j D_{p_j} D_{x_i} \psi + \frac{1}{2} \dot{\xi}^i \dot{\xi}^j D_{p_i} D_{p_j} \psi \right\} dt \\ &\equiv \frac{1}{2} D_{x_i} D_{x_j} \varphi(x_0^*) \xi_0^i \xi_0^j + \int_0^T \left\{ \frac{1}{2} \xi_t^i A_t^{ij} \xi_t^j + \xi_t^i B_t^{ij} \dot{\xi}_t^j + \frac{1}{2} \dot{\xi}_t^i q_t^{ij} \dot{\xi}_t^j \right\} dt \end{aligned} \quad (24)$$

where the matrix-valued functions of time A , B and q are defined by (16). This quadratic functional of ξ characterizes the (approximate) Gaussian distribution of the perturbation. Let us notice that because x^* is assumed to be a minimizer of the action functional, this quadratic functional of ξ must be non-negative.

Now suppose that we have some C^1 symmetric-matrix-valued function of time, θ , such that $\theta_T = 0$. Then we may write

$$\begin{aligned} Q(\xi) &= Q(\xi) + \frac{1}{2} \xi_0 \cdot \theta_0 \xi_0 + \left[\frac{1}{2} \xi_t \cdot \theta_t \xi_t \right]_0^T \\ &= \frac{1}{2} \xi_0 \cdot (D^2 \varphi(x_0^*) + \theta_0) \xi_0 \\ &\quad + \int_0^T \left\{ \frac{1}{2} \xi_t \cdot A_t \xi_t + \xi_t \cdot B_t \dot{\xi}_t + \frac{1}{2} \dot{\xi}_t \cdot q_t \dot{\xi}_t + \frac{1}{2} \xi_t \cdot \dot{\theta}_t \xi_t + \xi_t \cdot \theta_t \dot{\xi}_t \right\} dt. \end{aligned} \quad (25)$$

The quadratic form inside the integral can be expressed as

$$\frac{1}{2} \dot{\xi}_t \cdot q_t \dot{\xi}_t + \xi_t \cdot (B_t + \theta_t) \dot{\xi}_t + \frac{1}{2} \xi_t \cdot (A_t + \dot{\theta}_t) \xi_t = \frac{1}{2} (\dot{\xi}_t + K_t \xi_t) \cdot q_t (\dot{\xi}_t + K_t \xi_t), \quad (26)$$

where $K_t \equiv q_t^{-1}(B_t^T + \theta_t)$, provided

$$A_t + \dot{\theta}_t = K_t^T q_t K_t = (B_t + \theta_t) q_t^{-1} (B_t^T + \theta_t), \quad (27)$$

that is, provided θ solves the ODE (17). Therefore we have shown that the second statement of the theorem follows once we have established the first.

We now turn to the issue of the existence⁶ of a solution to (17). The proof that the quadratic ODE (17) has a solution exploits the only piece of information we have about the quadratic form Q , namely that it is *non-negative*, since x^* maximized $\Lambda(x|y)$. We may therefore consider the optimal control problem

$$G(t, x) \equiv \inf \left\{ \int_t^T \left\{ \frac{1}{2} \xi_s \cdot A_s \xi_s + \xi_s \cdot B_s \dot{\xi}_s + \frac{1}{2} \dot{\xi}_s \cdot q_s \dot{\xi}_s \right\} ds : \xi_t = x \right\} \quad (28)$$

⁶Uniqueness of the solution is not a problem, because the right-hand side of the ODE (17) is locally Lipschitz in θ , and once we know that there exists a solution this is enough to guarantee uniqueness.

which is well posed, and the solution G has a quadratic form:

$$G(t, x) = \frac{1}{2} x \cdot H_t x, \quad (29)$$

which is centred, since clearly $G(t, 0) = 0$ for all t . Since G is the value function, we shall have that

$$h(t) \equiv \int_0^t \left\{ \frac{1}{2} \xi_s \cdot A_s \xi_s + \xi_s \cdot B_s \dot{\xi}_s + \frac{1}{2} \dot{\xi}_s \cdot q_s \dot{\xi}_s \right\} ds + G(t, \xi_t) \quad (30)$$

is non-decreasing whatever ξ we choose, and will be constant if we choose the optimal ξ . Differentiating (30) gives us

$$\dot{h}_t = \frac{1}{2} \xi_t \cdot A_t \xi_t + \xi_t \cdot B_t \dot{\xi}_t + \frac{1}{2} \dot{\xi}_t \cdot q_t \dot{\xi}_t + \frac{1}{2} \xi_t \cdot \dot{H}_t \xi_t + \xi_t H_t \dot{\xi}_t. \quad (31)$$

Minimizing over $\dot{\xi}_t$ and equating to zero to find the optimal $\dot{\xi}_t$ gives us

$$\dot{\xi}_t = -q_t^{-1} (B_t^T \xi_t + H_t \xi_t), \quad (32)$$

so that the minimized value of \dot{h}_t becomes

$$\frac{1}{2} \xi_t \cdot A_t \xi_t + \frac{1}{2} \xi_t \cdot \dot{H}_t \xi_t - \frac{1}{2} (B_t^T \xi_t + H_t \xi_t) \cdot q_t^{-1} (B_t^T \xi_t + H_t \xi_t).$$

At optimality, this must be zero; so

$$0 = A_t + \dot{H}_t - (B_t + H_t) q_t^{-1} (B_t^T + H_t). \quad (33)$$

Evidently from the definition we have $G(T, \cdot) \equiv 0$, so that H solves (17) with the zero boundary condition at $t = T$. In other words, we have found θ , and $\theta = H$. Notice in particular that the equation (32) becomes $\dot{\xi}_t = -K_t \xi_t$. □

Numerical aspects.

Theorem 3 characterizes the behaviour of the path x around the local minimizer x^* , but in any given example, this is not the end of the story. The point is that when we seek x^* we resort to numerical means⁷ to solve the ODE (11) subject to the boundary conditions (10), (12), and it may be that the solution found is not the global minimizer. It could be that the solution obtained is only a local minimizer; it could be that the solution found is a only stationary point of $\Lambda(\cdot|y)$; or it could be that the numerical method has got into difficulties and what it returns is not a solution. An obvious approach to this problem is simply to try numerical solution of (17) back in time from initial value $\theta_T = 0$; however, if the ODE solver fails, usually the entire calculation stops, and no attempt to continue is possible. We can avoid this issue if we transform the first-order non-linear Riccati ODE into a second-order linear ODE. The following result gives all we need.

⁷All the calculations for this paper were performed in Scilab www.scilab.org which is a free and very versatile package able to perform a wide range of numerical and graphical tasks.

Proposition 1 *Suppose that x^* is a solution to (10), (11), (12), and that A , B and q are defined in terms of x^* by (16). Then:*

(i) *If the ODE (17) has a unique solution, then the $n \times n$ -matrix-valued function F_t defined as the unique solution to the first-order linear ODE*

$$\dot{F}_t = -K_t F_t \equiv -q_t^{-1}(B_t^T + \theta_t)F_t, \quad F_T = I \quad (34)$$

solves the second-order linear ODE

$$0 = (A_t - \dot{B}_t^T)F_t + (B_t - B_t^T - \dot{q}_t)\dot{F}_t - q_t\ddot{F}_t \quad (35)$$

with boundary conditions

$$F_T = I, \quad B_T^T F_T + q_T \dot{F}_T = 0. \quad (36)$$

(ii) *If the (unique) solution to (35) with boundary conditions (36) is non-singular for all time, then θ defined by*

$$B_t^T + \theta_t = -q_t \dot{F}_t F_t^{-1} \quad (37)$$

solves (17), and is the unique solution.

(iii) *If the (unique) solution to (35) with boundary conditions (36) fails to be non-singular for all time, then x^* cannot be a local minimizer of $\Lambda(\cdot|y)$.*

PROOF. (i) Suppose that θ is the unique solution to (17), and use this to define a function F using (34) with conditions (36) at T . Since (34) is a linear ODE with bounded continuous coefficients, it has a unique solution. Multiplying (34) on the left by q_t and differentiating leads to

$$\begin{aligned} -q\ddot{F} - \dot{q}\dot{F} &= (\dot{B}^T + \dot{\theta})F + (B^T + \theta)\dot{F} \\ &= (\dot{B}^T - A + (B + \theta)q^{-1}(B^T + \theta))F + (B^T + \theta)\dot{F} \\ &= (\dot{B}^T - A)F - (B + \theta)\dot{F} + (B^T + \theta)\dot{F} \\ &= (\dot{B}^T - A)F - (B - B^T)\dot{F} \end{aligned} \quad (38)$$

which is the ODE (35), as required. Here, we have used (17) at line (38).

(ii) Using (37) to define a function θ of time, we see from (36) that $\theta_T = 0$, so it remains to check that θ so defined satisfies (17). Multiplying on the right by F and differentiating (37) gives us

$$(\dot{B}^T + \dot{\theta})F + (B^T + \theta)\dot{F} = -\dot{q}\dot{F} - q\ddot{F} = (\dot{B}^T - A)F + (B^T - B)\dot{F} \quad (39)$$

leading to the equation

$$(A + \dot{\theta})F = -(B + \theta)\dot{F} = (B + \theta)q^{-1}(B^T + \theta)F. \quad (40)$$

Since F is assumed to be invertible for all t , we can clear out the common factor of F on each side and recover the ODE (17) for θ .

(iii) Suppose that x^* were a local minimizer of $\Lambda(\cdot|y)$; then by Theorem 2, the ODE (17) has a unique solution. We may therefore define a $n \times n$ matrix-valued function f_t of time to solve

$$\dot{f}_t = -K_t f_t, \quad f_T = I$$

where $K_t = q_t^{-1}(B_t^T + \theta_t)$. Notice that K is bounded on $[0, T]$, so we may also define the $n \times n$ matrix-valued function φ_t by

$$\dot{\varphi}_t = \varphi_t K_t, \quad \varphi_T = I.$$

Calculus shows that $\varphi_t f_t$ is constant, and therefore equal to I ; in particular, f_t is always invertible. According to the first statement of the Proposition, f solves the ODE (35) subject to the boundary conditions (36), which has a unique solution. But by hypothesis, this solution becomes singular at some time in $[0, T]$, which is a contradiction. The conclusion follows. \square

REMARKS. The way Proposition 1 gets used in practice is that we firstly identify what we think may be the action-minimizing path x^* and then we solve the second-order linear ODE (35) with the boundary conditions (36). There is never any problem with this ODE, since it is linear with bounded continuous coefficients. Having solved for F , we now calculate the determinant of F_t along the path of F . If this ever vanishes, it indicates that x^* was not a local minimizer of Q_0 .

5 An extension.

We have so far considered a situation where a diffusion X is observed through some continuous process Y of observations, and we are required to estimate X . Another quite natural situation we might want to consider would be where the hidden process X is some positive diffusion, which acts as the intensity process of some observed point process. So suppose that there is some non-negative diffusion x satisfying

$$dx_t = \sigma(t, x_t) dW_t + \mu(t, x_t) dt \tag{41}$$

which serves as the stochastic intensity of a counting process N . The times $0 < \tau_1 < \tau_2 < \dots < \tau_{N_T} < T$ of events are observed, and we have to filter x from these observations. In this case, the action integral (= -log-likelihood) to be minimised is just

$$\varphi(x_0) + \int_0^T \psi(s, x_s, \dot{x}_s) ds + \int_0^T x_s ds - \sum_{i=1}^n \log x_{\tau_i}, \tag{42}$$

where we have abbreviated $n \equiv N_T$, and ψ is as before

$$\psi(s, x, p) = (p - \mu(s, x))^2 / 2\sigma(s, x)^2. \quad (43)$$

Once again, we consider a small perturbation ξ away from the optimal x , and equate the first-order change to zero. Assuming for simplicity that σ and μ are independent of time gives us⁸

$$\begin{aligned} 0 &= \varphi'(x_0)\xi_0 + \int_0^T \{\psi_x \xi + \psi_p \dot{\xi} + \xi\} dt - \sum_{i=1}^n \frac{\xi(\tau_i)}{x(\tau_i)} \\ &= \varphi'(x_0)\xi_0 + \sum_{j=0}^n \int_{\tau_j}^{\tau_{j+1}} \{\psi_x \xi + \psi_p \dot{\xi} + \xi\} dt - \sum_{i=1}^n \frac{\xi(\tau_i)}{x(\tau_i)} \\ &= \varphi'(x_0)\xi_0 + \sum_{j=0}^n \left[\int_{\tau_j}^{\tau_{j+1}} \{\psi_x - \psi_{px} \dot{x} - \psi_{pp} \dot{p} + 1\} \xi dt + \left[\xi_t \psi_p(x_t, \dot{x}_t) \right]_{\tau_j}^{\tau_{j+1}} \right] - \sum_{i=1}^n \frac{\xi(\tau_i)}{x(\tau_i)}. \end{aligned}$$

In order for this to be zero whatever perturbation ξ is used, we have to have a number of conditions:

$$0 = \psi_x - \psi_{px} \dot{x} - \psi_{pp} \dot{p} + 1 \quad \text{in each interval } (\tau_j, \tau_{j+1}) \quad ; \quad (44)$$

$$0 = \varphi'(x_0) - \psi_p(x_0, p_0) \quad ; \quad (45)$$

$$0 = -\psi_p(x_{\tau_i}, p_{\tau_i+}) + \psi_p(x_{\tau_i}, p_{\tau_i-}) - x(\tau_i)^{-1} \quad ; \quad (46)$$

$$0 = \psi_p(x_T, p_T). \quad (47)$$

Thus the least-action path must be constructed piecewise in each of the intervals between observations.

Example 5. To illustrate this situation, we suppose that the positive diffusion x is a log-Brownian motion:

$$dx_t = x_t(\sigma dW_t + \mu dt). \quad (48)$$

Some calculations reveal that the ODE (44) is

$$0 = \frac{p^2 - x\dot{p} + \sigma^2 x^3}{\sigma^2 x^3}. \quad (49)$$

Notice that any solution of the ODE (49) must be convex while it is positive. The non-linear ODE (49) has various solutions, and can indeed be solved explicitly. However, the explicit solutions are explosive, so it turns out that to calculate the solution it is more effective to work numerically, which allows us to truncate the coefficients to prevent explosions. Once a

⁸We write $\tau_0 = 0, \tau_{n+1} = T$.

solution is found, the truncation is no longer necessary. Since $\psi_p = (p - \mu x)/\sigma^2 x^2$, the final condition (47) tells us that we must have $p_T = \mu x_T$, and (46) gives that

$$\Delta p_{\tau_i} \equiv p_{\tau_i+} - p_{\tau_i-} = -\sigma^2 x_{\tau_i}, \quad (50)$$

which determines the change of derivative over the observation points. Since the gradient drops at the observation points, this allows that the solution x found does not have to be convex globally, even though we have observed that it must be convex in each interval between observations.

Figure 5 illustrates the results of applying the method to this example. Despite the fact that there were only 49 points observed in the observation period, the general shape of the hidden intensity is well captured.

6 Relationship to particle filtering.

For the kind of problems discussed in this paper, the particle filtering⁹ methodology provides a natural approach. But some simple thought experiments highlight significant limitations of the methodology, and suggests ways in which the approach of this paper could in some circumstances supplement or replace conventional particle filtering.

Arguably one of the greatest strengths of particle filtering is that it evolves an estimate of the *posterior distribution* of the hidden variables; this guards most effectively against the common mistake of underestimating the error in parameter estimation. However, the methodology though simple is hard to apply effectively. One place where problems arise is in evolving the posterior distribution (represented by a finite collection of N point masses) forward to the next time point. The simplest version of particle filtering allows each point to make a jump according to the Markovian transition mechanism, and then reweights the new particles according to their likelihood given the new observation. One often finds¹⁰ that the likelihood of almost all (or indeed, *all*) the new particles is tiny; most of the new randomly-selected particles are miles away from the observation.

To understand the magnitude of the effect, imagine that there is a ball B of radius b hidden in the unit cube $C = [0, 1]^d$ in \mathbb{R}^d ; think of this ball as the set where $f(y|x) > \varepsilon$, where $f(\cdot|x)$ is the density of the new observation given the true state x of the hidden Markov process, and $\varepsilon > 0$ is some fixed threshold value of likelihood. If points are cast at random into C , how many on average will be needed before one hits the ball B ? The answer of course is just $b^{-d} V_d^{-1}$, where

$$V_d \equiv \frac{(2\pi)^{d/2} 2^{-d/2}}{\Gamma(1 + \frac{1}{2}d)}$$

⁹The signal processing literature refers to *Sequential Monte Carlo (SMC)* methods.

¹⁰This is particularly problematic when there is Gaussian observation error.

is the volume of the unit ball in \mathbb{R}^d . If we suppose that the radius is $b = 0.1$, representing a moderate amount of observational noise, Figure 6 plots the mean number $b^{-d}V_d^{-1}$ against dimension d , with the line where the mean number is one million plotted as a red dashed line across the figure. This crosses the plot at dimension roughly equal to 7. Thus if we want on average *one* particle to fall in the ball of radius 0.1 around the observation in dimension 7, we will require about one million attempts. Now of course we would not in practice just mindlessly fire particles completely at random into C ; various forms of importance sampling will hugely improve the success rate. But this is not quite as straightforward as it first appears. If we had a representation for the observation as

$$Y_t = X_t + \varepsilon_t \tag{51}$$

where X is the hidden state of the Markov process, and ε_t is a Gaussian observation error, then we would bias the jumps from the particles in the population at time $t - 1$ towards the new observation Y_t . However, the observation equation (51) is not how it is in most examples; more generally we have

$$Y_t = \Phi(X_t) + \varepsilon_t \tag{52}$$

for some very complicated function Φ , and in order to bias the jumps of the particles towards values of x which are likely relative to Y_t *we have to know what values of x make $\Phi(x)$ close to Y_t* . In other words, we have in effect to find the ML estimator of X_t given Y_t . Thus the particle-filtering methodology, which is wholeheartedly Bayesian in philosophy and execution, is often forced to resort to frequentist methodology in practice to help it overcome the curse of dimension. The least-action approach advanced here is such a ML method, and could be used to pick out a most likely path x^* given a path Y of observations; the second-order analysis would tell us how to simulate paths around the most likely path x^* so as to generate a sample of paths which would have a reasonably high likelihood.

7 Conclusions.

This paper offers an approach to estimating a hidden diffusion process observed either through other components of a joint diffusion, or through a point process whose intensity is the hidden diffusion. The approach is in effect a maximum-likelihood approach, but because of the continuous time parameter, calculus-of-variations techniques can be applied to identify the least-action (=maximum-likelihood) path. Moreover, by taking the calculus-of-variations analysis to second order, we are able to find not only what the most likely path is given the observations, but also approximately what is the distribution of the hidden path about the most likely path. One point which is worth emphasizing is that in contrast to other approaches, this methodology works without any assumption of symmetrizability of the diffusion, and requires only calculus in \mathbb{R}^d . The numerical methods involve shooting solutions of ODEs, so this can be technically quite delicate, especially if the time interval over which the solutions are to be

constructed is long. Accordingly, it cannot yet be claimed that this methodology will work well in all conceivable situations. However, it is a methodology which will in principle work in quite high dimension, and for this reason shows considerable promise.

A Appendix

PROOF OF THEOREM 1. If $z : [0, T] \rightarrow \mathbb{R}^d$ is C^1 , then let $z^{(n)}$ denote the piecewise-linear approximation which matches z at each multiple of $h = 2^{-n}$. Then¹¹ we have

$$\begin{aligned}
\lambda_n(x^{(n)}|y) &= -\frac{1}{2} \sum_{j=0}^{N-1} h \left| \sigma(jh, z_{jh}^{(n)})^{-1} \left(\frac{z_{jh+h}^{(n)} - z_{jh}^{(n)}}{h} - \mu(jh, z_{jh}^{(n)}) \right) \right|^2 - \varphi(x_0) \\
&= -\frac{1}{2} \int_0^T \left| \sigma(s_n, z_{s_n}^{(n)})^{-1} \left(\dot{z}_{s_n} - \mu(s_n, z_{s_n}^{(n)}) \right) \right|^2 ds - \varphi(x_0) \\
&\rightarrow -\frac{1}{2} \int_0^T \left| \sigma(s, z_s)^{-1} \left(\dot{z}_s - \mu(s, z_s) \right) \right|^2 ds - \varphi(x_0) \\
&= \Lambda(x|y)
\end{aligned}$$

by dominated convergence. Hence we have

$$\limsup_{n \rightarrow \infty} \inf_x \{ -\lambda_n(x|y) \} \leq \inf_x \{ -\Lambda(x|y) \}. \quad (53)$$

In the other direction, let $L_n = \inf_x \{ -\lambda_n(x|y) \}$, and suppose that $x^{(n)}$ is close to minimizing $-\lambda_n(\cdot|y)$ in the sense that

$$-\lambda_n(x^{(n)}|y) \leq L_n + 2^{-n}. \quad (54)$$

We then extend $x^{(n)}$ by piecewise-linear interpolation to create a path $z^{(n)}$ defined throughout $[0, T]$. We shall compare the discrete integral

$$\begin{aligned}
D &\equiv \frac{1}{2} \sum_{j=0}^{N-1} h \left| \sigma(jh, z_{jh}^{(n)})^{-1} \left(\frac{z_{jh+h}^{(n)} - z_{jh}^{(n)}}{h} - \mu(jh, z_{jh}^{(n)}) \right) \right|^2 \\
&= \frac{1}{2} \int_0^T \left| \sigma(s_n, z_{s_n}^{(n)})^{-1} \left(\dot{z}_{s_n} - \mu(s_n, z_{s_n}^{(n)}) \right) \right|^2 ds \\
&\equiv \int_0^T |a_s|^2 ds
\end{aligned} \quad (55)$$

with the continuous integral

$$\begin{aligned}
C &\equiv \frac{1}{2} \int_0^T \left| \sigma(s, z_s^{(n)})^{-1} \left(\dot{z}_s - \mu(s, z_s^{(n)}) \right) \right|^2 ds \\
&\equiv \int_0^T |b_s|^2 ds
\end{aligned} \quad (56)$$

¹¹Recall that $s_n = 2^{-n} \lceil 2^n s \rceil$.

and show that the two are very close for large n . Notice that the only difference between (55) and (56) is that the time index passes through the discrete values $j2^{-n}$ in the first, but moves continuously through $[0, T]$ in the second. The comparison goes via

$$\begin{aligned}
|D - C| &\leq \int_0^T ||a_s|^2 - |b_s|^2| ds \\
&= \int_0^T |(a_s - b_s) \cdot (a_s + b_s)| ds \\
&\leq \left(\int_0^T |a_s - b_s|^2 ds \right) \left(\int_0^T |a_s + b_s|^2 ds \right)
\end{aligned} \tag{57}$$

by showing that the first factor in (57) is small, and the second is bounded. Firstly we exploit (54) to give for some positive finite constants Γ and γ and $0 \leq t \leq u \leq T$ the bounds

$$\begin{aligned}
\Gamma &\geq \int_t^u \left| \sigma(s_n, z_{s_n}^{(n)})^{-1} \left(\dot{z}_{s_n}^{(n)} - \mu(s_n, z_{s_n}^{(n)}) \right) \right|^2 ds \\
&\geq \gamma \int_t^u \left| \dot{z}_{s_n}^{(n)} - \mu(s_n, z_{s_n}^{(n)}) \right|^2 ds \\
&\geq \gamma \int_t^u \left\{ \frac{1}{2} \left| \dot{z}_{s_n}^{(n)} \right|^2 - \left| \mu(s_n, z_{s_n}^{(n)}) \right|^2 \right\} ds.
\end{aligned}$$

Hence we have the inequality¹²

$$\int_t^u \left| \dot{z}_{s_n}^{(n)} \right|^2 ds \leq \Gamma \tag{58}$$

which yields the modulus of continuity

$$\begin{aligned}
|z_u^{(n)} - z_t^{(n)}|^2 &= \left| \int_t^u \dot{z}_{s_n}^{(n)} ds \right|^2 \\
&\leq \Gamma \sqrt{u - t}.
\end{aligned} \tag{59}$$

It is easy to see that the analogous inequality

$$\left| \int_t^u \dot{z}_s^{(n)} ds \right|^2 \leq \Gamma \sqrt{u - t}. \tag{60}$$

also holds. The boundedness of the second factor in (57) now follows. Using the modulus of continuity, it is now straightforward to show that the first factor in (57) can be made arbitrarily small by taking n arbitrarily large.

Hence given $\varepsilon > 0$, for large enough n we have

$$\inf_x \{ -\Lambda(x|y) \} \leq -\Lambda(x^{(n)}|y) \leq -\lambda_n(x^{(n)}|y) + \varepsilon \leq L_n + 2^{-n} + \varepsilon.$$

¹² Γ denotes as usual a positive finite constant whose value matters little, and may change from line to line.

We conclude that

$$\inf_x \{ -\Lambda(x|y) \} \leq \liminf_{n \rightarrow \infty} L_n$$

which combines with (53) to give the statement (8).

For the final statement of Theorem 1, we use the condition (7) along with the modulus of continuity proved above. This shows that the family $\{x_n\}$ is equicontinuous, so by the Arzela-Ascoli Theorem is relatively compact in the topology of $C[0, T]$. Any subsequential limit of the x_n converges to a path for which the value of $-\Lambda(x|y)$ is minimal; but by assumption there is only one such path, x^* , and so any subsequential limit of the x_n must be x^* , hence the x_n converge in the (uniform) topology of $C[0, T]$ to x^* .

□

References

- [1] C. Archambeau, D. Cornford, M. Opper, and J. Shawe-Taylor. Gaussian process approximations of stochastic differential equations. In N. D. Lawrence, A. Schwaighofer, and J. Q. Candela, editors, *Gaussian Processes in Practice*, pages 1–16. Journal of Machine Learning Research, 2007.
- [2] P. Fearnhead, O. Papaspiliopoulos, G.O. Roberts, and A.M. Stuart. Random-weight particle filtering of continuous time stochastic processes. *Journal of the Royal Statistical Society B.*, 72:497–512, 2010.
- [3] X. Mao. *Stochastic Differential Equations and Applications*. Horwood Publishing, Chichester, 2007.
- [4] B. Markussen. Laplace approximation of transition densities posed as Brownian expectations. *Stochastic Processes and their Applications*, 2009.

OU process X observed as $Y+z$, where z is OU

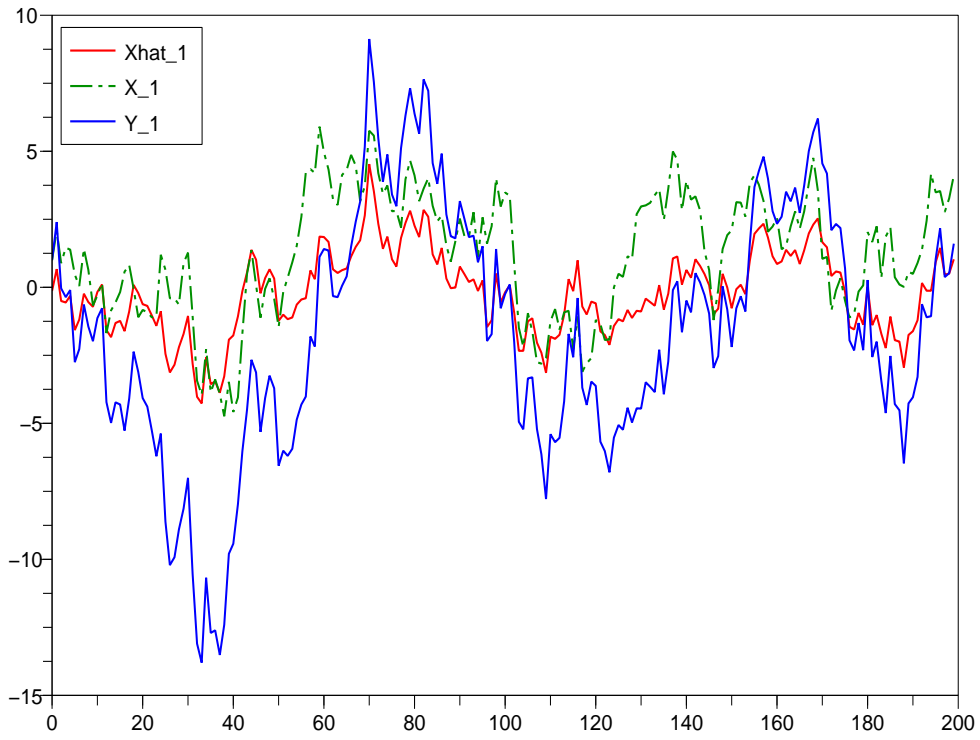


Figure 1: Result of least-action filtering.

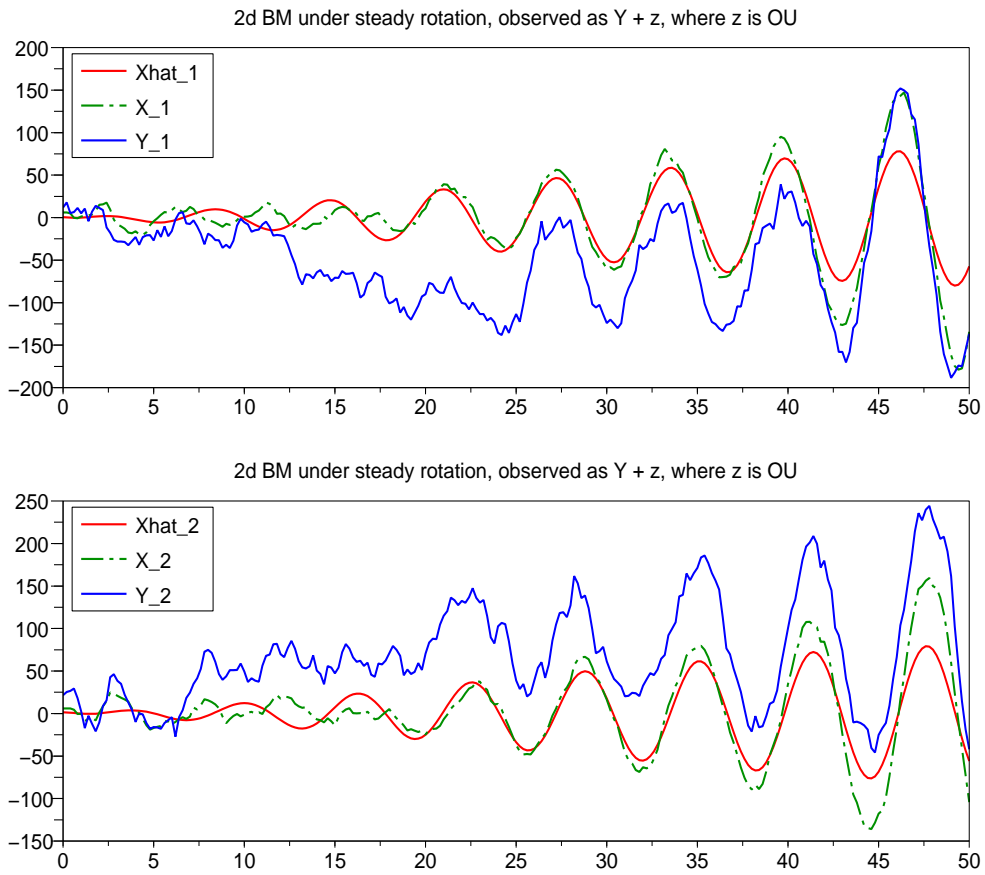


Figure 2: Result of least-action filtering.

$$dX = \text{sig}_X dW - b \sin(aX) dt, Y = X + \text{OU-noise}$$

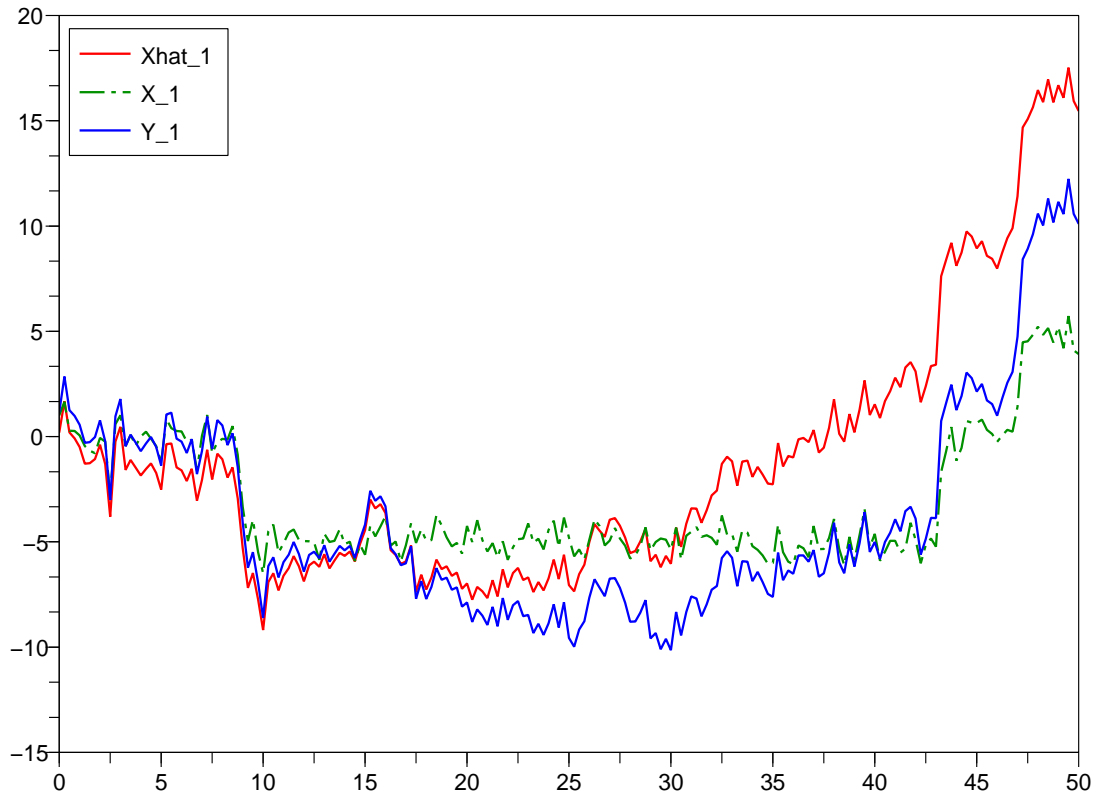


Figure 3: Result of least-action filtering.

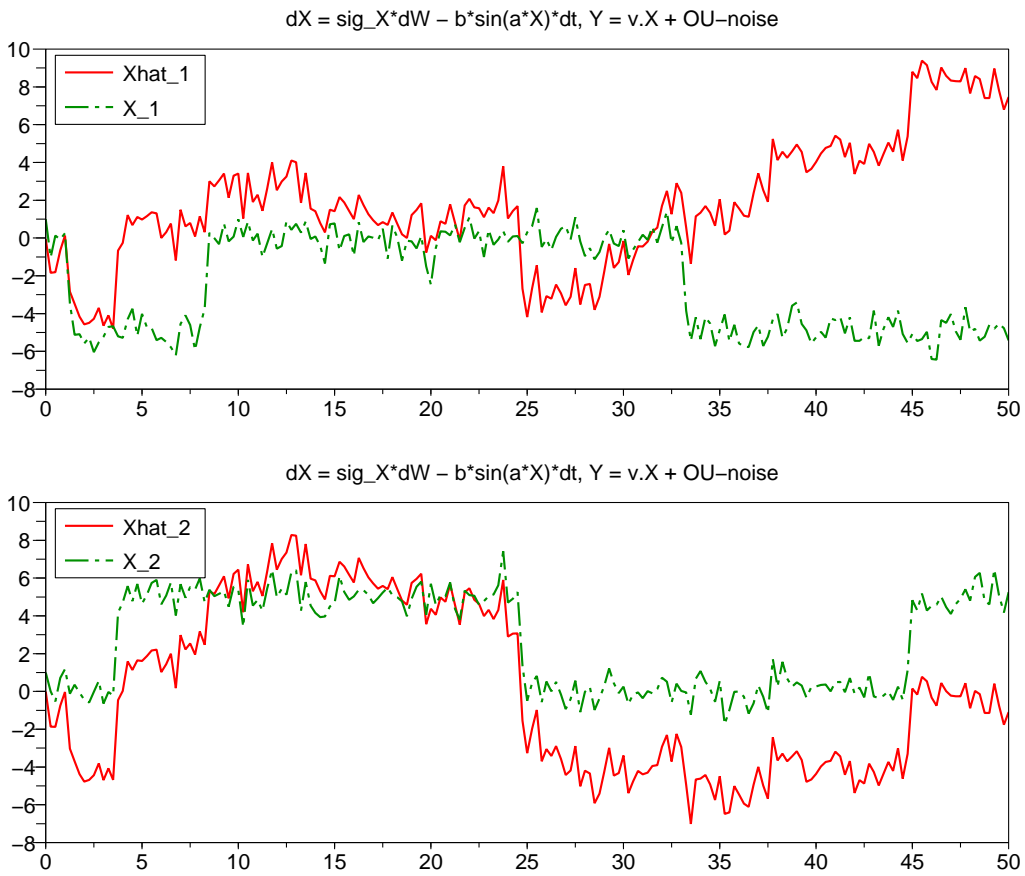


Figure 4: Result of least-action filtering.

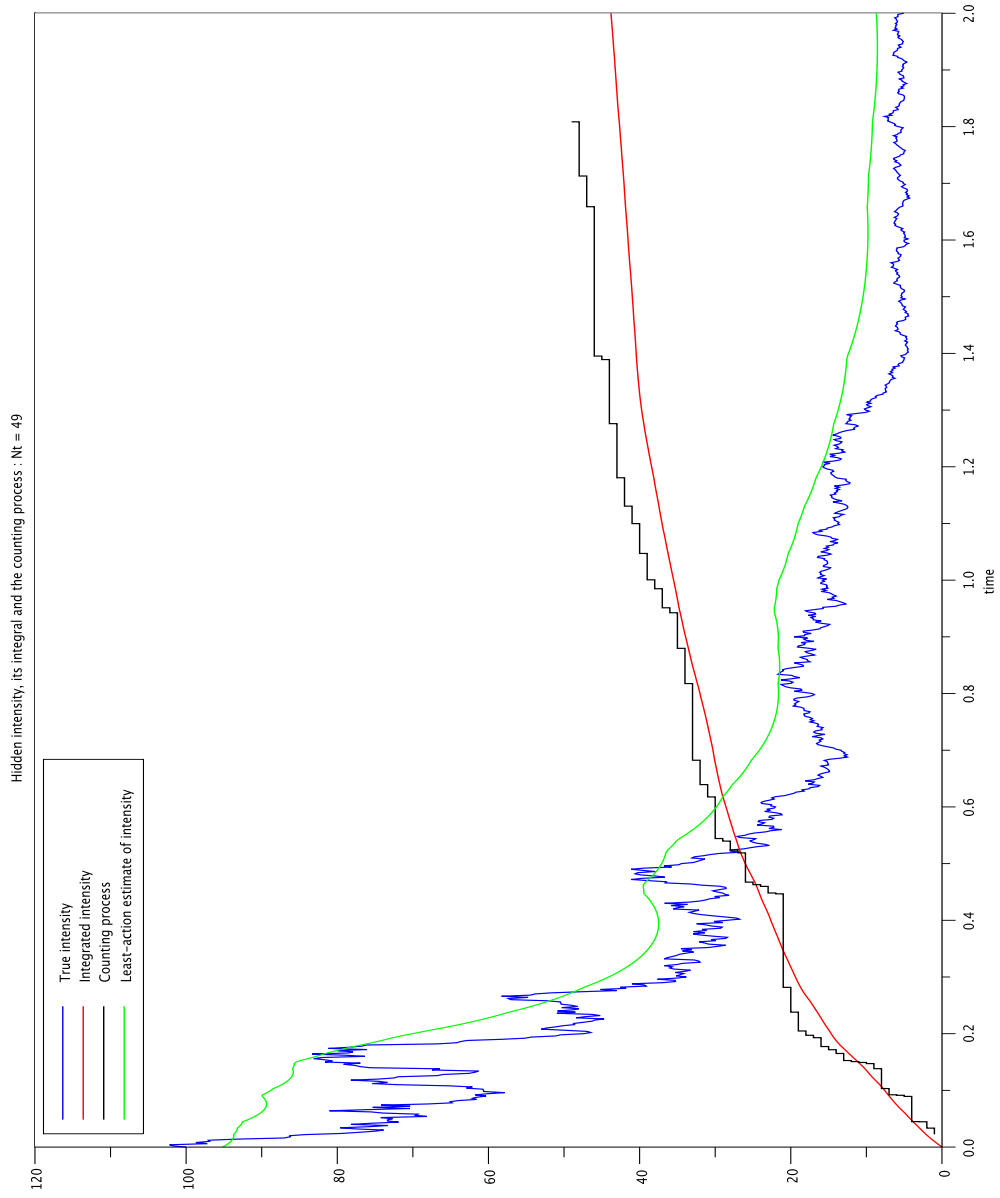


Figure 5: An example of least-action estimation of a hidden diffusion.

Log₁₀ of number of particles to get within 0.1 of hidden point

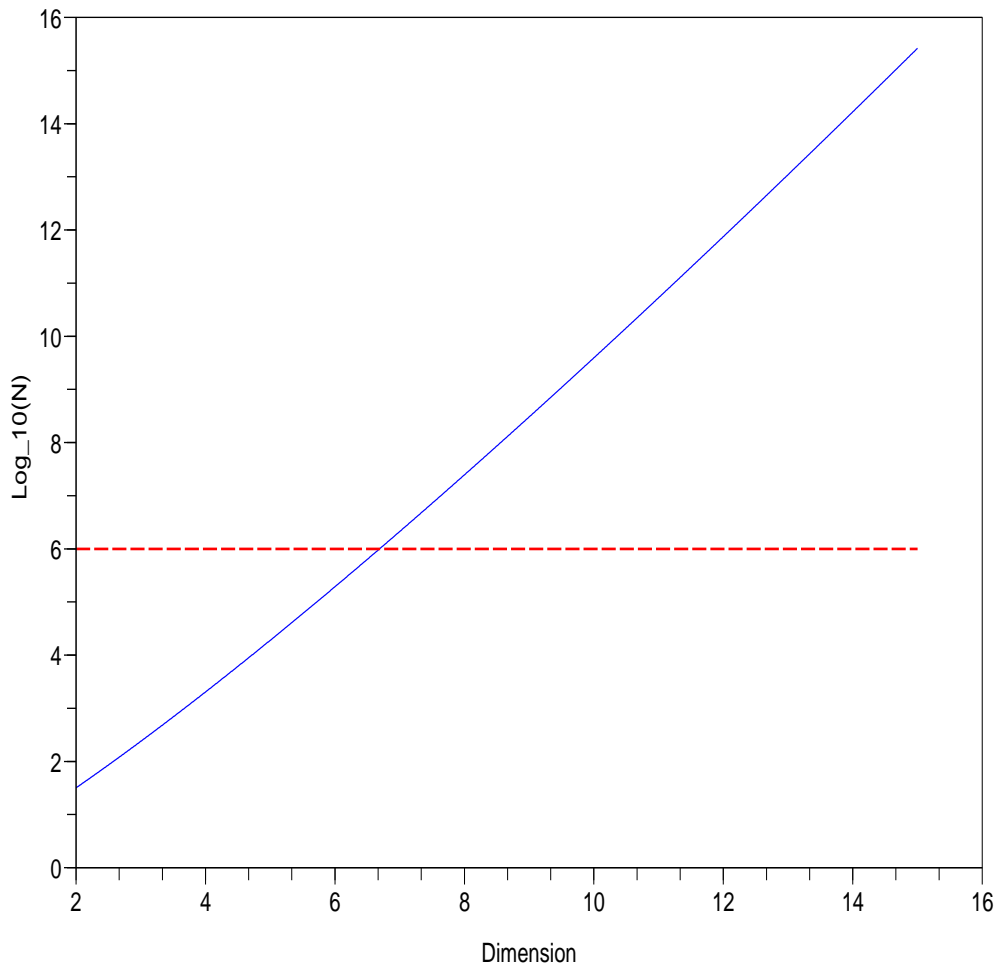


Figure 6: Limitations of particle filtering.