



Brief paper

Efficient grid-based Bayesian estimation of nonlinear low-dimensional systems with sparse non-Gaussian PDFs[☆]

Thomas R. Bewley^{a,1}, Atul S. Sharma^b^a Flow Control & Coordinated Robotics Labs, Dept of MAE, UC San Diego, La Jolla, CA 92093-0411, USA^b Department of Automatic Control and Systems Engineering, University of Sheffield, S1 3JD, UK

ARTICLE INFO

Article history:

Received 1 January 2010

Received in revised form

29 July 2011

Accepted 14 October 2011

Available online 3 May 2012

Keywords:

Nonlinear/non-Gaussian observer design

Grid-based Bayesian estimation

ABSTRACT

Bayesian estimation strategies represent the most fundamental formulation of the state estimation problem available, and apply readily to nonlinear systems with non-Gaussian uncertainties. The present paper introduces a novel method for implementing grid-based Bayesian estimation which largely sidesteps the severe computational expense that has prevented the widespread use of such methods. The method represents the evolution of the probability density function (PDF) in phase space, $p_{\mathbf{x}}(\mathbf{x}', t)$, discretized on a fixed Cartesian grid over *all* of phase space, and consists of two main steps: (i) between measurement times, $p_{\mathbf{x}}(\mathbf{x}', t)$ is evolved via numerical discretization of the Kolmogorov forward equation, using a Godunov method with second-order corner transport upwind correction and a total variation diminishing flux limiter; (ii) at measurement times, $p_{\mathbf{x}}(\mathbf{x}', t)$ is updated via Bayes' theorem. Computational economy is achieved by exploiting the localized nature of $p_{\mathbf{x}}(\mathbf{x}', t)$. An ordered list of cells with non-negligible probability, as well as their immediate neighbors, is created and updated, and the PDF evolution is tracked *only* on these active cells.

Published by Elsevier Ltd

1. Introduction

Bayesian estimation strategies are the most general class of solutions to the state estimation problem, and apply readily to nonlinear systems where information about the state is represented by a probability density function (PDF) of general form. In this paper, we introduce a novel, computationally cheap method for implementing grid-based Bayesian estimation that exploits the fact that the PDF is usually negligible in most of phase space, while avoiding many of the disadvantages of other methods. The idea of grid-based Bayesian estimation dates back at least to Stratonovich (1959, 1960). The equations underlying the algorithm are laid out clearly in Jazwinski (1970, p. 164), and are summarized below. However, numerical implementation of these equations has only been attempted sporadically in the half century since, for instance by Kramer and Sorenson (1988), Terwiesch and Agarwal (1994) and Ungarala, Chen, and Li (2006). Grid-based Bayesian methods typically suffer from the

twin burdens of high computational cost and a finite domain size; indeed, Arulampalam, Maskell, Gordon, and Clapp (2002), in their otherwise insightful review of particle filter methods, all but dismiss grid-based methods in Section IV.B. We believe that this level of pessimism on this class of methods is unwarranted. The algorithm developed, dubbed GBES (Grid-based Bayesian Estimation Exploiting Sparsity), provides a means of efficient computation by building on an accurate integration scheme for hyperbolic systems, and a novel gridding scheme over *all* of phase space.

2. Grid-based Bayesian estimation exploiting sparsity

Consider the state estimation of the nonlinear system

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, \mathbf{w}), \quad \mathbf{y} = \mathbf{h}(\mathbf{x}, \mathbf{v}). \quad (1)$$

The grid-based Bayesian estimation method is best visualized as an evolution of the PDF of the state estimate $\hat{\mathbf{x}}$ discretized on a fixed grid over all of phase space \mathbb{R}^n ; assuming the state \mathbf{x} develops according to the nonlinear equation (1), the method consists of two relatively straightforward steps (for details, see Jazwinski, 1970, p. 164):

- (i) Between measurement times, the PDF itself, $p_{\mathbf{x}}(\mathbf{x}', t)$, is marched via discretization of the Kolmogorov forward equation

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Giuseppe De Nicolao under the direction of Editor Ian R. Petersen.

E-mail addresses: bewley@ucsd.edu (T.R. Bewley), a.s.sharma@sheffield.ac.uk (A.S. Sharma).

¹ Tel.: +1 858 534 4287; fax: +1 858 822 3107.

(also called the *Fokker–Planck equation*)

$$\frac{\partial p_{\mathbf{x}}(\mathbf{x}', t)}{\partial t} = -\frac{\partial f_i(\mathbf{x}', t) p_{\mathbf{x}}(\mathbf{x}', t)}{\partial x'_i} + \frac{1}{2} \frac{\partial^2 q_{ij} p_{\mathbf{x}}(\mathbf{x}', t)}{\partial x'_i \partial x'_j}, \quad (2)$$

where summation over repeated indices is implied and q_{ij} is the (i, j) th element of the spectral density, Q , of the state disturbances (note that, in the special case that Q is diagonal and the state disturbance is independent of \mathbf{x} , this is just a diffusion term). Risken (2002) reviews a number of methods for solving this equation, including analytic methods for special cases and eigenfunction expansions (focusing on the stationary solution), however state-space and time discretization for a non-stationary solution (as performed here) is only mentioned briefly, and assumed to apply only to finite domains. An accurate numerical method for marching this equation in time in the case that $Q = 0$ is outlined in Section 3; adding an appropriate term to this discretization to apply diffusion to the PDF (to account for Gaussian state disturbances) is straightforward, as discussed in Section 3.2.

- (ii) At the measurement times t_k , the PDF is updated via Bayes' theorem (Bayes, 1763),

$$p_{\mathbf{x}}(\mathbf{x}', t_{k+}) = \frac{p_{\mathbf{y}}(\mathbf{y}_k | \mathbf{x}') p_{\mathbf{x}}(\mathbf{x}', t_{k-})}{C}, \quad (3)$$

where $p_{\mathbf{x}}(\mathbf{x}', t_{k+})$ denotes the *a posteriori* PDF (after accounting for the measurement \mathbf{y}_k), $p_{\mathbf{y}}(\mathbf{y}_k | \mathbf{x}')$ denotes the uncertainty associated with the measurement (which may or may not be Gaussian in \mathbf{x}'), $p_{\mathbf{x}}(\mathbf{x}', t_{k-})$ denotes the *a priori* PDF (before accounting for the measurement \mathbf{y}_k), and C is an appropriate normalization constant, which is selected for every measurement update to normalize the discretization of $p_{\mathbf{x}}(\mathbf{x}', t_{k+})$ such that its integral over phase space is unity.

To understand how the continuous state-space is discretized, recall that the cumulative distribution function (CDF) of a random real vector \mathbf{x} , denoted $f_{\mathbf{x}}(\mathbf{x})$, maps $\mathbf{x} \in \mathbb{R}^n$ to the real interval $[0, 1]$ that monotonically increases in each of the components of \mathbf{x} , and is defined

$$f_{\mathbf{x}}(\mathbf{x}) = P(x_1 \leq \underline{x}_1, x_2 \leq \underline{x}_2, \dots, x_n \leq \underline{x}_n),$$

where $\underline{\mathbf{x}}$ is some particular value of the random vector \mathbf{x} and $P(S)$ denotes a probability measure that the conditions stated in S are true. For any random vector \mathbf{x} whose CDF is differentiable everywhere, the probability density function (PDF) $p_{\mathbf{x}}(\mathbf{x}') \geq 0$ is a scalar function of \mathbf{x}' defined such that

$$f_{\mathbf{x}}(\mathbf{x}) = \int_{-\infty}^{\underline{x}_1} \int_{-\infty}^{\underline{x}_2} \dots \int_{-\infty}^{\underline{x}_n} p_{\mathbf{x}}(\mathbf{x}') dx'_1 dx'_2 \dots dx'_n,$$

$$\Leftrightarrow p_{\mathbf{x}}(\mathbf{x}') = \left. \frac{\partial^n f_{\mathbf{x}}(\mathbf{x})}{\partial x_1 \partial x_2 \dots \partial x_n} \right|_{\mathbf{x}=\mathbf{x}'}$$

For small $|\Delta \mathbf{x}'|$, the quantity $p_{\mathbf{x}}(\mathbf{x}') \Delta x'_1 \Delta x'_2 \dots \Delta x'_n$ represents the probability that the random vector \mathbf{x} takes some value within a small rectangular region centered at the particular value \mathbf{x}' and of width $\Delta x'_i$ in each coordinate direction \mathbf{e}_i .

The method we have developed maintains a list of active cells on the grid over all of phase space in order to limit both the computational effort and the memory storage required in the numerical simulation. This list includes all cells in the discretization for which the PDF is greater than a given threshold, as well as all cells which, though they may or may not themselves exceed this threshold, are either one of the two immediate neighbor cells, in each of the n coordinate directions, of those cells which exceed the threshold, or are one of the four neighbor cells, in each of the ${}_n C_2$ pairs of coordinate directions, which touch a corner of the cells which exceed the threshold. At each time

step, cells are added to and removed from this list as appropriate, and the fluxes initialized and updated on every interior boundary between adjacent cells in the list. When performing a computation restricted to an evolving list of active grid cells of this sort, the relative position of the various cells in the list is needed frequently. This may be determined efficiently by keeping in each list record a pointer to the two immediate neighbor cells in each coordinate direction in the list, if these neighbor cells are present in the list, or to a null record if not, and updating these pointers appropriately as records are added to and removed from the list.² These pointers interconnecting the list facilitate rapid computation of the numerical discretization given in (5) in the next section. The most expensive step in maintaining this list of neighbor cells is making the appropriate connections when a new record is added to the list. Though this may be accomplished by scanning the entire list, this approach becomes prohibitively expensive as the length of the list grows to thousands of cells. Instead, we keep the list ordered by its indices (e.g., in a phase space with $n = 3$, ordered first by i , then by j , then by k), and store the elements of the list as a binary tree. This allows the time-limiting search step to proceed at $\mathcal{O}(N \log N)$ operations, where N is the number of list elements. Conveniently, this list ordering and searching can be handled using the C++ Standard Template Library map container.

3. Accurate numerical integration of the Kolmogorov forward equation

The PDE governing the evolution of the PDF in the present problem is given by (2). If $Q = 0$, the equation is hyperbolic; if $Q > 0$, the equation, strictly speaking, changes type to elliptic. In practice, however, Q is usually relatively small. It is thus fitting to design a numerical method for accurate simulation of (2) based on a proven algorithm for accurate simulation of hyperbolic PDEs. Fortunately, the fluid mechanics community has focused on the development of high performance computing techniques for numerical simulation of such “convection-dominated” problems for over 40 years, and these techniques are now quite refined and well understood. The numerical method best suited to the present problem is somewhat involved; a comprehensive review of this class of methods is given in LeVeque (2002). To focus this discussion, consider first the two-dimensional, linear, hyperbolic PDE in conservation form

$$\frac{\partial p(x, y, t)}{\partial t} = -\frac{\partial u(x, y) p(x, y, t)}{\partial x} - \frac{\partial v(x, y) p(x, y, t)}{\partial y}, \quad (4)$$

noticing that higher-dimensional cases follow as an obvious extension. Following Chapters 4, 6, 9, 19, and 20 of LeVeque (2002), we implement a Godunov-type finite volume method by writing (4) on a uniform Cartesian 2D mesh (with constant Δx and Δy) in the form

$$\frac{p_{ij}^{n+1} - p_{ij}^n}{\Delta t} = -\frac{F_{i+1/2,j}^n - F_{i-1/2,j}^n}{\Delta x} - \frac{G_{i,j+1/2}^n - G_{i,j-1/2}^n}{\Delta y}, \quad (5)$$

where the fluxes $F_{i-1/2,j}^n$ and $G_{i,j-1/2}^n$ are determined, for all i and j , by first initializing

$$F_{i-1/2,j}^n = u_{i-1/2,j}^+ p_{i-1,j}^n + u_{i-1/2,j}^- p_{i,j}^n,$$

$$G_{i,j-1/2}^n = v_{i,j-1/2}^+ p_{i,j-1}^n + v_{i,j-1/2}^- p_{i,j}^n,$$

² The four neighbor cells, in each pair of coordinate directions, which touch the corner of a given cell may be found by referencing the neighbor cell of a neighbor cell.

where $u^+ = \max(u, 0)$, $u^- = \min(u, 0)$, etc., then applying the *corner transport upwind* (CTU) terms by updating, for all i and j ,

$$\begin{aligned} F_{i-1/2,j-1}^n &\leftarrow F_{i-1/2,j-1}^n - \Delta t \frac{u_{i-1/2,j-1}^- v_{i,j-1/2}^-}{2} \frac{\Delta p_{i,j-1/2}^n}{\Delta y}, \\ F_{i+1/2,j-1}^n &\leftarrow F_{i+1/2,j-1}^n - \Delta t \frac{u_{i+1/2,j-1}^+ v_{i,j-1/2}^-}{2} \frac{\Delta p_{i,j-1/2}^n}{\Delta y}, \\ F_{i-1/2,j}^n &\leftarrow F_{i-1/2,j}^n - \Delta t \frac{u_{i-1/2,j}^- v_{i,j-1/2}^+}{2} \frac{\Delta p_{i,j-1/2}^n}{\Delta y}, \\ F_{i+1/2,j}^n &\leftarrow F_{i+1/2,j}^n - \Delta t \frac{u_{i+1/2,j}^+ v_{i,j-1/2}^+}{2} \frac{\Delta p_{i,j-1/2}^n}{\Delta y}, \\ G_{i-1,j-1/2}^n &\leftarrow G_{i-1,j-1/2}^n - \Delta t \frac{v_{i-1,j-1/2}^- u_{i-1/2,j}^-}{2} \frac{\Delta p_{i-1/2,j}^n}{\Delta x}, \\ G_{i-1,j+1/2}^n &\leftarrow G_{i-1,j+1/2}^n - \Delta t \frac{v_{i-1,j+1/2}^+ u_{i-1/2,j}^-}{2} \frac{\Delta p_{i-1/2,j}^n}{\Delta x}, \\ G_{i,j-1/2}^n &\leftarrow G_{i,j-1/2}^n - \Delta t \frac{v_{i,j-1/2}^- u_{i-1/2,j}^+}{2} \frac{\Delta p_{i-1/2,j}^n}{\Delta x}, \\ G_{i,j+1/2}^n &\leftarrow G_{i,j+1/2}^n - \Delta t \frac{v_{i,j+1/2}^+ u_{i-1/2,j}^+}{2} \frac{\Delta p_{i-1/2,j}^n}{\Delta x}, \end{aligned}$$

where $\Delta p_{i-1/2,j}^n = p_{ij}^n - p_{i-1,j}^n$, $\Delta p_{i,j-1/2}^n = p_{ij}^n - p_{i,j-1}^n$, then applying the *high-resolution correction terms* by updating, for all i and j ,

$$\begin{aligned} F_{i-1/2,j}^n &\leftarrow F_{i-1/2,j}^n + \Delta t \frac{|u_{i-1/2,j}|}{2} \left(\frac{\Delta x}{\Delta t} - |u_{i-1/2,j}| \right) \\ &\quad \times \frac{\Delta p_{i-1/2,j}^n}{\Delta x} \phi(\theta_{i-1/2,j}^n), \\ G_{i,j-1/2}^n &\leftarrow G_{i,j-1/2}^n + \Delta t \frac{|v_{i,j-1/2}|}{2} \left(\frac{\Delta y}{\Delta t} - |v_{i,j-1/2}| \right) \\ &\quad \times \frac{\Delta p_{i,j-1/2}^n}{\Delta y} \phi(\theta_{i,j-1/2}^n), \end{aligned}$$

where

$$\theta_{i-1/2,j}^n = \begin{cases} \Delta p_{i-3/2,j}^n / \Delta p_{i-1/2,j}^n & \text{if } u_{i-1/2,j} \geq 0, \\ \Delta p_{i+1/2,j}^n / \Delta p_{i-1/2,j}^n & \text{if } u_{i-1/2,j} < 0, \end{cases}$$

$$\theta_{i,j-1/2}^n = \begin{cases} \Delta p_{i,j-3/2}^n / \Delta p_{i,j-1/2}^n & \text{if } v_{i,j-1/2} \geq 0, \\ \Delta p_{i,j+1/2}^n / \Delta p_{i,j-1/2}^n & \text{if } v_{i,j-1/2} < 0, \end{cases}$$

and the *flux limiter function* $\phi(\theta) \in [0, 2]$ is selected as one of several possible choices, including the *monotonized central-difference* (MC) limiter and the *van Leer* limiter:

$$\text{MC} : \phi(\theta) = \max\{0, \min[(1 + \theta)/2, 2, 2\theta]\},$$

$$\text{van Leer} : \phi(\theta) = (\theta + |\theta|)/(1 + |\theta|).$$

Note that exact conservation of the discrete approximation of the integral of p over phase space, as implied by the continuous formulation in (4), follows immediately from (5).

3.1. Numerical analysis

In regions characterized by smooth variation of p , $\theta \approx 1$ and $\phi(\theta) \approx 1$, and the algorithm described in Section 3 is amenable to straightforward numerical analysis. For simplicity, consider here the 1D test problem

$$\frac{\partial p}{\partial t} = -u \frac{\partial p}{\partial x} \quad (6)$$

where u is a positive or negative constant. In this case, the discretization described above reduces to

$$\frac{p_i^{n+1} - p_i^n}{\Delta t} = - \frac{F_{i+1/2}^n - F_{i-1/2}^n}{\Delta x}$$

where

$$F_{i-1/2}^n = \frac{u}{2} (p_i^n + p_{i-1}^n) - \frac{u^2 \Delta t}{2 \Delta x} (p_i^n - p_{i-1}^n),$$

and thus

$$\frac{p_i^{n+1} - p_i^n}{\Delta t} = -u \frac{(p_{i+1}^n - p_{i-1}^n)}{2 \Delta x} + \frac{u^2 \Delta t}{2} \frac{(p_{i+1}^n - 2p_i^n + p_{i-1}^n)}{(\Delta x)^2}.$$

Now applying to this equation the multidimensional Taylor series expansion,

$$\begin{aligned} y_{i+k}^{n+m} &= y_i^n + m \Delta t \left(\frac{\partial y}{\partial t} \right)_i + k \Delta x \left(\frac{\partial y}{\partial x} \right)_i + \frac{(m \Delta t)^2}{2} \left(\frac{\partial^2 y}{\partial t^2} \right)_i \\ &\quad + \frac{(k \Delta x)^2}{2} \left(\frac{\partial^2 y}{\partial x^2} \right)_i + m \Delta t k \Delta x \left(\frac{\partial^2 y}{\partial x \partial t} \right)_i + \dots, \end{aligned}$$

and rearranging appropriately, gives

$$\begin{aligned} (p_t)_i^n &= -u (p_x)_i^n - \frac{\Delta t}{2} (p_{tt})_i^n + \frac{u^2 \Delta t}{2} (p_{xx})_i^n \\ &\quad + \mathcal{O}((\Delta t)^2, (\Delta x)^2, \Delta x \Delta t). \end{aligned}$$

Differentiating (6) with respect to t and inserting (6) into the RHS of the result, it is seen that the second and third terms on the RHS of the above expression cancel. Thus, in regions of smooth variation of p , the proposed scheme is *second-order accurate in both space and time*.³ A similar analysis follows for problems in higher dimensions.

3.2. Accounting for diffusion

A diffusion term is easily added to the discretization given in (5) in a second-order central finite difference fashion simply by updating the fluxes such that, for all i and j ,

$$\begin{aligned} F_{i-1/2,j}^n &\leftarrow F_{i-1/2,j}^n + \bar{\mu} \frac{\Delta p_{i-1/2,j}^n}{\Delta x}, \\ G_{i,j-1/2}^n &\leftarrow G_{i,j-1/2}^n + \bar{\mu} \frac{\Delta p_{i,j-1/2}^n}{\Delta y}, \end{aligned}$$

where $\bar{\mu}$ is the coefficient of the diffusion term that is applied numerically. The flux limiter functions mentioned at the end of Section 3 are designed to reduce the algorithm, locally, to a first-order spatial behavior while applying sufficient numerical diffusion in regions of large local curvature of p on the grid, to provide a total variation diminishing (TVD) solution (that is, preventing spurious oscillations with new local minima and maxima). We may compensate for the diffusion introduced by the numerical discretization of the convective terms simply by appropriately reducing the diffusion $\bar{\mu}$ applied in the numerical simulation of (2).

3.3. Validation

A simple yet sensitive numerical test of the algorithm is given in Fig. 1; this numerical test was taken with $u = y$ and $v = -x$ in order to give simple solid body rotation about the origin, as suggested by LeVeque (2002). If $Q = 0$, the exact solution of the test problem considered in Fig. 1 at $t = 2\pi$, after a single rotation of the system about the origin, is simply the initial condition. For the case $Q = 2\mu l$ where μ is constant and positive, the exact solution of this problem at $t = 2\pi$ may be obtained analytically by means of Fourier transforms. As seen by comparing Fig. 1 to Figure 20.5 of LeVeque (2002), the result obtained via the GBEES approach is essentially identical to that obtained using the complete

³ Meaning that the error is bounded by a term proportional to $(\Delta x)^2$ in space and $(\Delta t)^2$ in time, giving convergence of $\mathcal{O}(1/N^2)$.

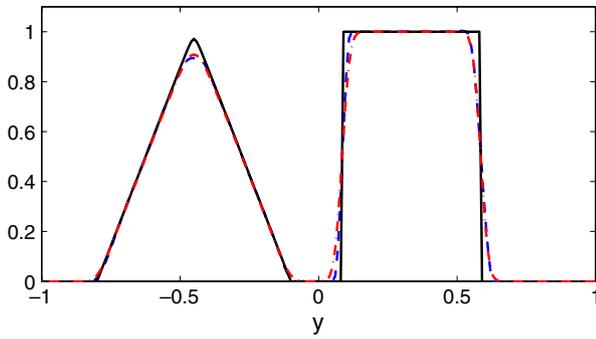


Fig. 1. Computation of a sensitive test system (governed by (4) with $u(x, y) = y$ and $v(x, y) = -x$), using the method in Section 3, taking $Q = 0$, $\Delta x = 0.01$, $\Delta t = 0.001$, and a threshold of 10^{-3} (that is, tracking numerically only those cells with $p > 0.001$ and their immediate neighbors). The distributions are compared in cross section in the plane $x = 0$, with the black solid line denoting the exact solution at $t = 2\pi$ for $Q = 0$, the blue dashed line denoting the numerical solution at $t = 2\pi$ using $Q = 0$, $\Delta x = 0.01$, $\Delta t = 0.001$ and the red dot-dashed line denoting the exact solution at $t = 2\pi$ for $Q = 8 \times 10^{-5}$. The leading-order error in the numerical solution at $t = 2\pi$ is a small amount of diffusion of p . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

grid when sufficiently small threshold, time step and state-space discretization is used. The information loss due to the discretization scheme may be quantified via the Kullback–Liebler divergence (Kullback & Liebler, 1951), $D_{KL}(P_0, P_\epsilon)$, where a distribution P_ϵ is used to approximate the true distribution P_0 . The Kullback–Liebler divergence using the simulation in Fig. 1 when compared to the true analytic solution is 0.089 bits (with the distributions normalized to integrate to unity), whereas the divergence from the true case to the diffusion case calculated analytically for $Q = 8 \times 10^{-5}$ is 0.085 bits.⁴ The divergence from the true solution to a simulation using a truncation threshold of 10^{-16} (not shown) is almost the same as the more aggressively truncated example and is visually indistinguishable. As evident by comparing the numerical solution at $t = 2\pi$ in Fig. 1 to the initial condition, the discretization described in Section 3 introduces a small numerical error in regions of high curvature. However, by comparing the numerical solution, for $\mu = 0$, $\Delta y = 0.001$ and $\Delta x = 0.01$, to the exact solution, for $\mu = 4 \times 10^{-5}$, it is evident that the leading-order error of the numerical discretization is just a bit of additional diffusion, the level of which may be determined by a suitable minimization process.

4. Numerical results

A Bayesian approach is justified when the uncertainty of the estimate is significantly non-Gaussian, such as in the estimation of a nonlinear system with relatively large uncertainty, leaving us with particle filtering or grid-based methods; what is perhaps still uncertain is the numerical tractability of a grid-based approach when one exploits the sparsity of the PDF in the manner described in Section 2. Thus, in order to test the efficiency of the GBEES algorithm, as well as to demonstrate how it can capture with unprecedented accuracy the evolution of a non-Gaussian PDF, we have applied the GBEES algorithm to the estimation of the three-state Lorenz system

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}), \quad \mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}, \quad \mathbf{f}(\mathbf{x}) = \begin{pmatrix} \sigma(x_2 - x_1) \\ -x_2 - x_1x_3 \\ -bx_3 + x_1x_2 - br \end{pmatrix}$$

⁴ This value of Q was chosen so that the divergence from the true case to the diffusion case was close to that of the true case to the numerical solution.

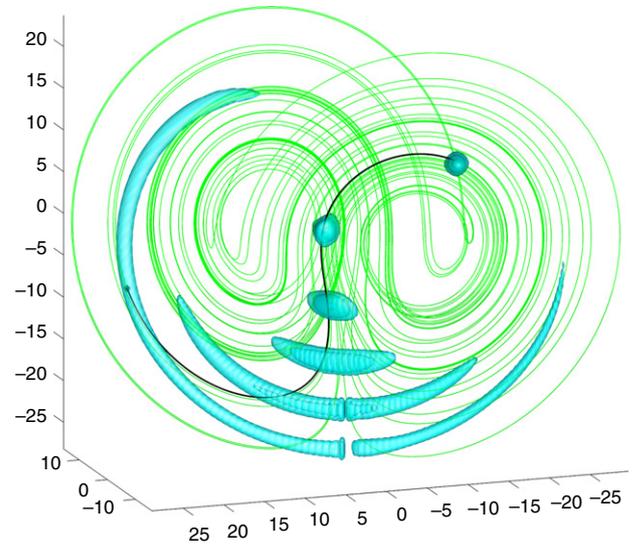


Fig. 2. The rapid transformation of a PDF from Gaussian to highly non-Gaussian in the Lorenz system, with no measurement updates. Visualized are $p = 0.005$, $p = 0.0005$, and $p = 0.00005$ isosurfaces of the PDF in phase space at $t = 0$ (in the upper-right), $t = 0.2$, $t = 0.4$, $t = 0.6$, $t = 0.8$, and $t = 1$. The simulation was performed with a variable time step $\Delta t \leq 0.001$, a grid spacing of $\Delta x = 0.25$, and a threshold of $\epsilon = 10^{-6}$. This simulation required less than 40 s of computation on a 2009 vintage Apple laptop computer (2.4 GHz Intel Core 2 Duo) using a single-threaded C++ implementation of the present algorithm tracking about 50,000 active cells (of 5.3×10^6 total cells in the domain shown) at $t = 1$. The modest memory requirements are proportional to the number of cells.

with $\sigma = 4$, $b = 1$, and $r = 48$. For these parameter values, the system is chaotic, and the attractor takes the familiar form indicated by the green line in Fig. 2. Also illustrated in this Fig. 2 is the evolution of an initially (at $t = 0$) Gaussian PDF $p_{\mathbf{x}}(\mathbf{x}', t)$, the evolution of which is governed by the Kolmogorov equation (2), with no measurement updates applied and no added process noise (diffusion). The distribution narrows significantly in the direction normal to the attractor, and spreads out rapidly in the direction of the maximum local Lyapunov exponent along the attractor; by $t = 1$, the PDF is highly non-Gaussian. Note also in the $t = 0.8$ and $t = 1$ isosurfaces the remarkable division of the PDF into two distinct lobes in the vicinity of the x_3 axis (the vertical coordinate axis in the figures), which is invariant and unstable in the Lorenz system.

Fig. 3 represents the evolution of the PDF when measurements (with Gaussian uncertainty) of x_3 are taken at every time step. Computationally, the problem addressed in the figure is significantly easier than the “open-loop” problem illustrated in Fig. 2, as the number of active cells by $t = 1$ is reduced from 50,000 to only 4000; the computation time for this simulation is also reduced accordingly, from 40 to 4 s for the time interval shown. The PDF at time $t = 1$ splits into two concentrated regions on the left and right sides of the figure. Future measurements might reveal in which region the state really is; until such measurements are received, the GBEES algorithm is perfectly capable of following both. A plain Kalman filter, which assumes a central estimate, would necessarily fail to model such a splitting.

5. Analysis and conclusions

A novel algorithm is introduced in this paper to exploit the remarkable sparsity of the evolving PDF in phase space, thereby, for the first time, making high-resolution grid-based Bayesian estimation computationally tractable for nontrivial systems. The method generalizes straightforwardly to any number dimensions, with computational cost expected to be a trade-off between

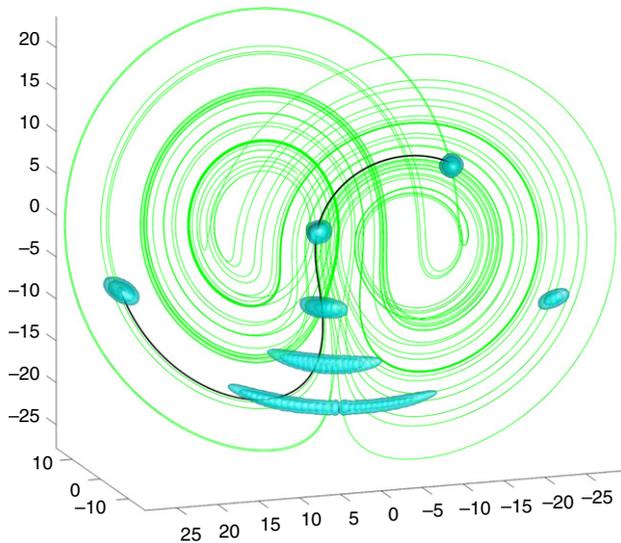


Fig. 3. The same simulation as Fig. 2 but with measurements of x_3 (vertical axis), with Gaussian uncertainty, at every time step. The black line is the 'true' state which generates the measurements. Note that, by $t = 1$, the PDF splits into two concentrated regions. This simulation required about 4 s of computation with the same hardware, tracking about 4000 active cells (of 5.3×10^6 total cells in the domain shown) at $t = 1$.

the curse of dimensionality and the increased sparseness of the PDF. In application, the algorithm developed is shown to track, with unprecedented fidelity, the completely non-Gaussian PDF of the estimate of a Lorenz system, both with and without measurement updates. The simulation exhibits a competition between information loss due to the random state disturbances and stretching of the PDF in the unstable directions of the system, and information gain from measurements.

Grid-based Bayesian estimation algorithms are sometimes referred to as *approximate* grid-based methods. We point out that the numerical analysis of Section 3.1 establishes that the numerical method used to propagate the Kolmogorov equation in the present grid-based estimation algorithm is *second-order accurate in both space and time*; this compares favorably to the (*slower than linear*) $\mathcal{O}(1/\sqrt{N})$ convergence rate of particle methods applied to the Kolmogorov equation (see Bernard, Talay, & Tubaro, 1994).

Finally, Lagrangian (that is, particle-based) simulation techniques have been explored for decades in the field of fluid mechanics, but for $n > 2$ remain mostly a research novelty. On the other hand, grid-based methods (often with adaptive grids to focus the computational effort where it is needed) have proven immensely successful in a variety of complex situations in fluid mechanics, such as in the characterization of fluid turbulence and in the design of commercial airliners, where computational methods have largely supplanted repetitive wind-tunnel testing. There appears to be no reason why the same success of grid-based methods will not also be realized in Bayesian estimation approaches,

once the community working on such problems fully appreciate how the remarkable sparsity of the PDF in such problems may be exploited.

Acknowledgments

The authors gratefully acknowledge Prof. Paulo Luchini for insightful discussions related to this work. An Imperial College Junior Research Fellowship (AS) is also gratefully acknowledged.

References

- Arulampalam, S., Maskell, S., Gordon, N., & Clapp, T. (2002). A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50, 174–188.
- Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, 370–418.
- Bernard, P., Talay, D., & Tubaro, L. (1994). Rate of convergence of a stochastic particle method for the Kolmogorov equation with variable coefficients. *Mathematics of Computation*, 63, 555–587.
- Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. Academic Press.
- Kramer, S. C., & Sorenson, H. W. (1988). Recursive Bayesian estimation using piecewise constant approximations. *Automatica*, 24, 789–801.
- Kullback, S., & Liebler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22, 79–86.
- LeVeque, R. J. (2002). *Finite volume methods for hyperbolic problems*. Cambridge University Press.
- Risken, H. (2002). *The Fokker–Planck equation*. Springer.
- Stratonovich, R. L. (1959). Optimum nonlinear systems which bring about a separation of a signal with constant parameters from noise. *Radiofizika*, 2, 892–901.
- Stratonovich, R. L. (1960). Application of the Markov processes theory to optimal filtering. *Radio Engineering and Electronic Physics*, 5, 1–19.
- Terwiesch, P., & Agarwal, M. (1994). A discretized nonlinear state estimator for batch processes. *Computers & Chemical Engineering*, 19, 155–169.
- Ungarala, S., Chen, Z., & Li, K. (2006). Bayesian state estimation of nonlinear systems using approximate aggregate Markov chains. *Industrial & Engineering Chemistry Research*, 45, 4208–4221.



Thomas R. Bewley received a B.S./M.S. from Caltech in 1989, a diploma from the von Karman Institute for Fluid Mechanics in 1990, and a Ph.D. from Stanford in 1998. He is now a Professor in the MAE Department at UC San Diego, where he directs the UCSD Flow Control and Coordinated Robotics Labs. Research interests include the control, estimation, forecasting, and adaptive observation of aerodynamic and environmental fluid systems, and the design, optimization, and control of small multimodal robotic vehicles.



Atul S. Sharma received the MSci. degree in physics from University College London, UK, in 1998 and the Ph.D. degree in tokamak modeling and control from Imperial College, London, UK, in 2002. He was a Research Associate and then a Junior Research Fellow at Imperial College before being appointed to his current position as Lecturer at the University of Sheffield. His interests include the modeling and control of turbulent fluid flows.