HAL open science

Particle observers for contracting dynamical systems

Silvere Bonnabel, Jean-Jacques Slotine

To cite this version:

Silvere Bonnabel, Jean-Jacques Slotine. Particle observers for contracting dynamical systems. 3rd conference on Geometric Science of Information (GSI'17), Nov 2017, Paris, France. hal-01692404

HAL Id: hal-01692404

https://minesparis-psl.hal.science/hal-01692404

Submitted on 25 Jan 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Particle observers for contracting dynamical systems

Silvère Bonnabel[1] and Jean-Jacques Slotine[2]

[1] MINES ParisTech, PSL Reasearch University, Centre for Robotics, 60 bd Saint-Michel, 75006 Paris, France, silvere.bonnabel@mines-paristech.fr,
[2] Department of Mechanical Engineering, Massachusetts Institute of Technology, MA, 02139 USA, jjs@mit.edu.

1 Introduction

In the present paper we consider a class of partially observed dynamical systems. As in the Rao-Blackwellized particle filter (RBPF) paradigm (see e.g., [Doucet et al. (2000)]), we assume the state x can be broken into two sets of variables $x = (z, r)$ and has the property that conditionally on z the system's dynamics possess geometrical contraction properties, or is amenable to such a system by using a nonlinear observer whose dynamics possess contraction properties. Inspired by the RBPF we propose to use particles to approximate the r variable and to use a simple copy of the dynamics (or an observer) to estimate the rest of the state. This has the benefits of 1- reducing the computational burden (a particle filter would sample the variable x also), which is akin to the interest of the RBPF, 2- coming with some indication of stability stemming from contraction (actual proofs of stability seem difficult), and 3- the obtained filter is well suited to systems where the dynamics of x conditionally on z is precisely known and the dynamics governing the evolution of z is quite uncertain.

2 A primer on contraction theory

2.1 Background on contraction theory

Consider a Riemannian manifold $(\mathcal{M}, g)$, where g denotes the metric. Consider local coordinates. In the present paper, we will simplify the exposure by systematically assuming that $\mathcal{M} = \mathbb{R}^n$. The squared infinitesimal length is given by the quadratic form:

$$\|dx\|^2 = \sum_{1 \leq i,j \leq n} g_{ij}(x) dx_i dx_j$$

The matrix $G = (g_{ij})_{1 \leq i,j \leq n}$ is called the Riemannian metric tensor and it generally depends on x. Now, consider the continuous time deterministic system described by the following ordinary differential equation (ODE) on $\mathbb{R}^n$:

$$\frac{d}{dt} x = f(x), \tag{1}$$

with f a smooth nonlinear function satisfying the usual conditions for global existence and unicity of the solution. For a detailed proof of the following theorem, see e.g., [Pham and Slotine (2013)].

Theorem 1 ([Lohmiller and Slotine(1998)]) *Let $J_f(x)$ denote the Jacobian matrix of $f(x)$. Assume that $M(x) = G^T(x)G(x)$ is uniformly positive definite, and that $G(x)J_f(x)G^{-1}(x)$ is uniformly negative definite, then all trajectories exponentially converge to a single trajectory. Moreover, the convergence rate is equal to $\lambda > 0$ which is the supremum over x of the largest eigenvalue of $G(x)J_f(x)G^{-1}(x)$. More precisely, if $a(t)$ and $b(t)$ are two trajectories of* (1), *we have:*

$$d_g(a(t), b(t)) \leq d_g(a(0), b(0))e^{-2\lambda t},$$

where d_g denotes the Riemannian distance associated to metric g.

2.2 Nonlinear observers for contracting systems

Consider the system (1) where $x(t) \in \mathbb{R}^N$, with partial observations

$$y(t) = h(x(t)), \tag{2}$$

The goal of observer design, is to estimate in real time the **unknown** quantity $x(t)$ with the greatest possible accuracy given all the measurements up to current time t. Assume that for a class of functions $y(t)$, the dynamics

$$\frac{d}{dt}z = f(z) + K(z, y)(y - h(z)) \tag{3}$$

can be proved to be contractive with rate $\lambda > 0$. Then, the observer for the system (1)-(2) defined by

$$\frac{d}{dt}\hat{x} = f(\hat{x}) + K(\hat{x}, y)(y - h(\hat{x})), \tag{4}$$

possesses convergence properties. Indeed, as the simulated $\hat{x}(t)$ and the true trajectory $x(t)$ are both solutions of equation (3), Theorem 1 applies and we have:

$$d_g(\hat{x}(t), x(t)) \leq d_g(\hat{x}(0), x(0))e^{-2\lambda t}.$$

3 The basic particle observer

3.1 The Rao-Blackwellized particle filter (RBPF)

Consider a (discrete) Markov process r_t of initial distribution $p(r_0)$ and transition equation $p(r_t \mid r_{t-1})$. The variable r_t is hidden, and assume we have as observation a random variable y_t at time t, which is correlated with r_t. The observations are assumed to be conditionally independent given the process r_t.

The goal of discrete time filtering is to infer online the hidden variables from the observations, that is, to compute:

$$p(r_t \mid y_{1:t}), \quad \text{where} \quad y_{1:t} = \{y_1, \cdots, y_t\},$$

or more generally $p(r_{1:t} \mid y_{1:t})$. Assume now, that we also want to infer another related process z_t, such that $p(z_t \mid y_{1:t}, r_{1:t})$ can be analytically evaluated. This is typically the case using a Kalman filter when conditionally on r the system is linear and Gaussian. A simple version of the RBPF is given by Algorithm 1.

Algorithm 1 RBPF with prior sampling (see e.g., [Doucet et al. (2000)])

Draw N particles from the prior initial distribution $p(r_0)$

loop

Sample from the prior

$$r_t^{(i)} \sim p(r_t \mid r_{t-1}^{(i)}), \quad \text{and let} \quad r_{1:t}^{(i)} = \left(r_t^{(i)}, r_{1:t-1}^{(i)}\right)$$

Evaluate and update weights

$$w_t^{(i)} = p(y_t \mid y_{1:t-1}, r_{1:t}^{(i)})\ w_{t-1}^{(i)}$$

Normalize weights

$$\tilde{w}_t^{(i)} = \frac{w_t^{(i)}}{[\sum_j w_t^{(j)}]^{-1}}$$

The estimate of the expected value $\mathbb{E}\left(F(z,r)\right)$ of any function F is

$$\sum_i \tilde{w}_t^{(i)} \mathbb{E}_{p(z_t \mid y_{1:t}, r_{1:t}^{(i)})} \left(F(r_t^{(i)}, z_t)\right)$$

Resample if necessary, i.e., duplicate and suppress particles to obtain N random samples with equal weights (i.e., equal to $1/N$).

end loop

3.2 The particle observer for conditionally contracting systems

Consider a noisy dynamical system of the form

$$\frac{d}{dt} z = f(z, r) \tag{5}$$

$$\frac{d}{dt} r = g(r) + w(t) \tag{6}$$

where f, g are smooth maps, $w(t)$ is a process noise, and we have an initial prior distribution $\pi_0(z, r)$ at time $t = 0$. Assume one has access to discrete time uncertain measurements $y_n = h(z_{t_n}, r_{t_n}) + V_n$ at times $t_0 < t_1 < t_2 < \cdots$, and where V_n are unknown independent identically distributed random variables with

known density l, that is, $p(y \mid z, r) = l(y - h(z, r))$. We introduce the following definition.

Definition 1. *The system* (5)-(6) *is said to be a contraction conditionally on* z *if equation* (5) *is a contraction when* $r(t)$ *is considered as a known input.*

The rationale of our particle observer is as follows. If $r(t)$ were known, then, all trajectories of the system (5) would converge to each other due to the conditional contraction properties we assume. Thus, if we call $\hat{z}(t)$ a solution of (5) associated to some trajectory $\{r(t)\}_{t\geq 0}$, then asymptotically we have $p(z(t) \mid \{r(s)\}_{0\leq s\leq t}) \approx \delta(z(t) - \hat{z}(t))$, which means that contrarily to the RBPF paradigm we can not compute the conditional distributions in closed form but we have access to relevant approximations to them. Thus, letting $(\hat{z}_t^{(i)}, r_t^{(i)})$ be a solution to the stochastic differential equations (5)-(6), we have the following approximations that stem from the partial contraction properties of the system:

$$p(z_t \mid y_{1:t}, r_{1:t_n}^{(i)}) \approx \delta(z_t - \hat{z}_t^{(i)}), \quad p(y_{t_n} \mid y_{1:t_{n-1}}, r_{1:t_n}^{(i)}) \approx l(y_{t_n} - h(\hat{z}_{t_n}^{(i)}, r_{t_n}^{(i)})).$$

Thus, resorting to those approximation, and applying the RBPF methodology to the above system (5)-(6) we propose the following Algorithm 2.

Algorithm 2 The PO with prior sampling

Draw N particles $(z_{t_0}^{(1)}, r_{t_0}^{(1)}), \cdots, (z_{t_0}^{(N)}, r_{t_0}^{(N)})$ from the prior initial distribution $\pi_0(z, r)$

loop

Sample $(z_{t_n}^{(i)}, r_{t_n}^{(i)})$ from the prior by numerically integrating the stochastic differential equations (5)-(6) from time t_{n-1} to t_n.

Evaluate and update weights

$$w_t^{(i)} = l(y_{t_n} - h(z_{t_n}^{(i)}, r_{t_n}^{(i)}))\ w_{t-1}^{(i)}.$$

Numerically enforce that at least one weight is not equal to zero.

Normalize weights

$$\tilde{w}_t^{(i)} = \frac{w_t^{(i)}}{[\sum_j w_t^{(j)}]^{-1}}.$$

The estimate of the expected value $\mathbb{E}\left(F(z, r)\right)$ of any function F is approximated by

$$\sum_i \tilde{w}_t^{(i)} F(r_t^{(i)}, z_t^{(i)})$$

In particular the state is approximated by

$$\sum_i \tilde{w}_t^{(i)} (r_t^{(i)}, z_t^{(i)})$$

Resample if necessary, i.e., duplicate and suppress particles to obtain N random samples with equal weights (i.e., equal to $1/N$).

end loop

3.3 Some comments on the choice of model

The relevance of system (5)-(6) is debatable, for the two following reasons. First, it might be surprising that the dynamics of z conditionally on r be deterministic, whereas the dynamics of r be noisy. Second, because it is rare to find systems that are naturally (conditionally) contracting. Both issues will be partly addressed in the extensions outlined in the sequel. At this stage, we can make the following comments regarding the first issue. Assume both equations (5)-(6) to be noisy. Then, thanks to the contraction property, the asymptotic distribution of $z(t)$ conditionally on $r(t)$ is not very dispersed if the process noise is moderate, see [Pham et al. (2009)]. So the method may yield good results in practice. Assume on the other hand, both equations (5)-(6) to be deterministic. Then, it is hopeless to estimate and track efficiently the state with a (RB) particle filter, as the state space will not be explored adequately. Indeed, because multiple copies are produced after each resampling step, the diversity of the particle system decreases to a few points, which can be very different from the true state. To solve this degeneracy problem, the regularized particle filter was proposed in [Musso and Oujdane (1998)]. Albeit debatable, this technique may yield good results in practice. Following this route, we can postulate noisy equation (6) to implement our particle filter.

Remark 1 *Note that, here we do* not *deal with parameter identification, as in e.g., [Saccomani et. al (2003)]. Although this might look similar, $r(t)$ is not a parameter, preventing us to directly apply the results of e.g., [Wills et. al (2008)]*

4 A chemical reactor example

4.1 Retained model

Consider the exothermic chemical reactor of [Adebekun and Schork(1989)]. It was shown in [Lohmiller and Slotine(1998)] that, if the temperature T is known, and thus can be considered as an input, then the system is a contraction. But to achieve best performance, and filter the noise out of the temperature measurements, the temperature should be considered as a (measured) part of the state as in [Adebekun and Schork(1989)]. This leads to a system that is not a contraction. To make our point, we even propose to slightly modify the temperature dynamics to make it clearly unstable, yielding the more challenging following system:

$$\frac{d}{dt}I = \frac{q(t)}{V}(I_f - I) - k_d e^{-\frac{E_d}{RT(t)}} I \tag{7}$$

$$\frac{d}{dt}M = \frac{q(t)}{V}(M_f - M) - 2k_p e^{-\frac{E_d}{RT(t)}} M^2 I \tag{8}$$

$$\frac{d}{dt}P = \frac{q(t)}{V}(P_f - P) + k_p e^{-\frac{E_d}{RT(t)}} M^2 I \tag{9}$$

$$\frac{d}{dt}T = \beta T + \sigma_2 w(t) \tag{10}$$

where $w(t)$ is a white Gaussian standard noise, and $\sigma_2 > 0$ a parameter encoding the noise amplitude. Letting $V_n \sim \mathcal{N}(0,1)$ a random standard centered Gaussian, we assume discrete temperature measurements of the form:

$$y_n = T(t_n) + \sigma_1 V_n. \tag{11}$$

[Lohmiller and Slotine(1998)] already proved the system is contracting conditionally on $T(t)$. Thus, we can use the method described in Algorithm 2.

4.2 Simulation results

The true system is simulated according to the equations (7)-(8)-(9)-(10) where we turned the noise off in equation (10) (this means we started from a noise-free system for which the RBPF would not work properly, and used the regularization technique discussed at Section 3.3). The noisy output (11) was also simulated, where an observation is made every 5 steps. We chose $\tilde{\sigma}_2 = 0.1$. Density l is dictated by the observation noise, that is, $l(u) = \frac{1}{\sigma_1\sqrt{2\pi}} \exp(-\frac{u^2}{2\sigma_1^2})$ with $\sigma_1 = 1K$.

To apply our methodology, we assume that we have plausible physical upper bounds on the concentrations inside, and denote them by $I_{\max}, M_{\max}, P_{\max}$ and we let π_0 be the uniform distribution on the hyperrectangle $[0, I_{\max}] \times [0, M_{\max}] \times [0, P_{\max}]$ with a Dirac on the measured initial temperature. In the simulation, all those upper bounds are set equal to $4mol$. $\frac{q(t)}{V}$ and $T(t)$ are slowly oscillating around $1min^{-1}$ and $300K$, we have $k_d e^{+\frac{E_d}{RT(t)}} \approx 0.8min^{-1}$ and $k_p e^{+\frac{E_d}{RT(t)}} \approx 0.2L\ mol^{-1}\ min^{-1}$. We also let $\beta = 0.01min^{-1}$.

N=15 particles are used (which results in a very cheap to implement particle filter, as each particle is associated only to a naive observer). We resampled[3] each time the number of effective particles $1/(\sum_1^M w_j^2)$ drops below $N/4$, i.e., 25% of the total population. The resampling step is *necessary*, so that all particles gradually improve their estimation of the temperature, allowing the concentrations to be well estimated in turn.

The noise is efficiently filtered and all values asymptotically very well recovered, although a very reduced number of particles is used (15 observers are running in parallel) and measured temperature is noisy. See Figures 1 and 2.

5 Possible extensions and concluding remarks

Possible extensions are twofold. First, if equation (5) is noisy, one can use the same RBPO. Using the result of [Pham et al. (2009),Pham and Slotine (2013)], we can have an approximation of the asymptotic variance associated to the

[3] i.e., draw N particles from the current particle set with probabilities proportional to their weights; replace the current particle set with this new one. Instead of setting the weights of the new particles equal to $1/N$ as in the standard methodology, we preferred in the simulations to assign them their former weight and then normalize.

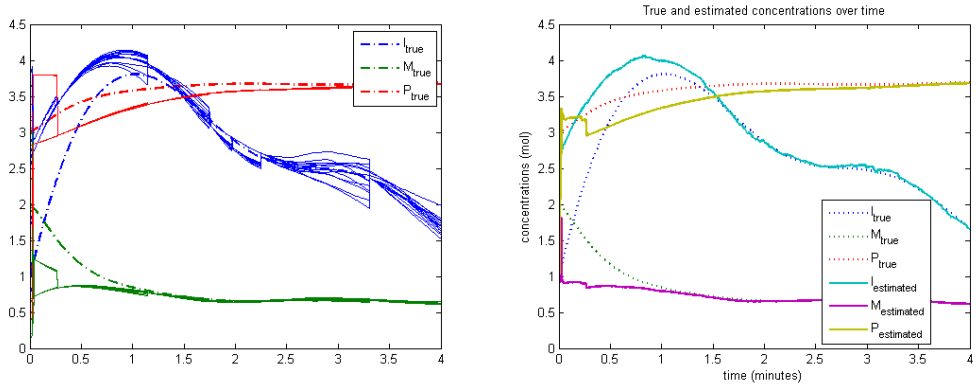

Fig. 1. Left: True concentrations (dashed lines) and trajectories of the 15 particles. We see the effect of resampling, that refocuses the bundle of trajectories on the fittest ones, when too many become unlikely. Right: True concentrations (dashed lines), and estimates of the particle observer (solid lines).

distribution $p(z(t) \mid \{r(s)\}_{0 \leq s \leq t})$. Thus, a Gaussian approximation to this distribution can be leveraged to implement a RBPF. Second, if f is not contracting conditionally on r, *but*, is amenable to it using an observer of the form

$$\frac{d}{dt} z = f(z) + K(z, y, r)(y - h(z))$$

then the method may still be applied.

The ideas introduced in this short paper might also be applied to differentially positive systems [Forni and Sepulchre(2016),Bonnabel et al.(2011)]. In the future, we would also like to study the behavior of particle filters for systems with contraction properties. A starting point could be to seek how to use the recent results of [Pham et al. (2009),Pham and Slotine (2013),Tabareau et. al (2010)] on stochastic contraction.

References

[Adebekun and Schork(1989)] Adebekun, D.K. and Schork, F.J. (1989). Continuous solution polymerization reactor control. 1. nonlinear reference control of methyl methacrylate polymerization. *Industrial & engineering chemistry research*, 28(9), 1308–1324.

[Aminzarey and Sontag(2014)] Aminzarey, Z. and Sontag, E.D. (2014). Contraction methods for nonlinear systems: A brief introduction and some open problems. In *53rd IEEE Conference on Decision and Control*, 3835–3847. IEEE.

[Bonnabel et al.(2011)] Bonnabel, S., Astolfi, A., and Sepulchre, R. (2011). Contraction and observer design on cones. In *IEEE Conference on Decision and Control 2011 (CDC11)*.

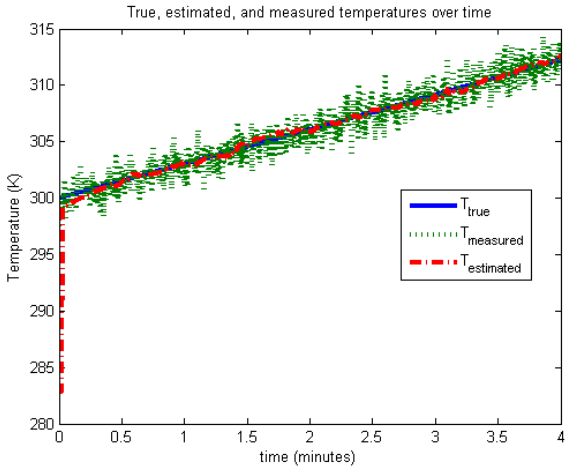

Fig. 2. True (solid line), measured (noisy line), estimated (dashed line, output by the RBPO) temperatures.

[Doucet et al. (2000)] Doucet, A. , De Freitas, N. and Murphy, K. (2000). Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, 176–183.

[Forni and Sepulchre(2016)] Forni, F. and Sepulchre, R. (2016). Differentially positive systems. *IEEE Transactions on Automatic Control*, 61(2), 346–359.

[Lohmiller and Slotine(1998)] Lohmiller, W. and Slotine, J. (1998). On contraction analysis for nonlinear systems. *Automatica*, 34(6), 683–696.

[Musso and Oujdane (1998)] C. Musso and N. Oudjane (1998). Regularization schemes for branching particle systems as a numerical solving method of the nonlinear filtering problem. *Proceedings of the Irish Signals and Systems Conference.*

[Pham et al. (2009)] Pham, Q.C., Tabareau, N., and Slotine, J.J. (2009). A contraction theory approach to stochastic incremental stability. *IEEE Transactions on Automatic Control*, 54(4), 816–820.

[Pham and Slotine (2013)] Pham, Q. C. and Slotine, J. J. (2013). Stochastic contraction in Riemannian metrics. *arXiv preprint arXiv:1304.0340.*.

[Saccomani et. al (2003)] Saccomani, M. P., Audoly, S. and D'Angi, L. (2003). Parameter identifiability of nonlinear systems: the role of initial conditions. *Automatica, 39(4), 619-632.*

[Tabareau et. al (2010)] Tabareau, N., Slotine, J. J. and Pham, Q. C. (2010). How synchronization protects from noise. *PLoS Comput Biol, 6(1), e1000637.*

[Wang and Slotine(2005)] Wang, W. and Slotine, J.J.E. (2005). On partial contraction analysis for coupled nonlinear oscillators. *Biological cybernetics*, 92(1), 38–53.

[Wills et. al (2008)] Wills, A., Schn, T. B and Ninness, B. (2008). Parameter estimation for discrete-time nonlinear systems using EM. *IFAC Proceedings Volumes, 41(2), 4012-4017.*