On Sequential Learning for Parameter Estimation in Particle Algorithms for State-Space Models

Chunlin Ji[1]

[1] Kuang-Chi Institute of Advanced Technology, Shenzhen, China

Correspondence: Chunlin Ji, Kuang-Chi Institute of Advanced Technology, Shenzhen 518000, China. Tel: 86-755-8632-9077. E-mail: chunlin.ji@kuang-chi.org

Abstract

Particle methods, also known as Sequential Monte Carlo, have been ubiquitous for Bayesian inference for state-space models, particulary when dealing with nonlinear non-Gaussian scenarios. However, in many practical situations, the state-space model contains unknown model parameters that need to be estimated simultaneously with the state. In this paper, We discuss a sequential analysis for combined parameter and state estimation. An online learning method is proposed to approach the distribution of the model parameter by tuning a flexible proposal mixture distribution to minimize their Kullback-Leibler divergence. We derive the sequential learning method by using a truncated Dirichlet processes normal mixture and present a general algorithm under a framework of the auxiliary particle filtering. The proposed algorithm is verified in a blind deconvolution problem, which is a typical state-space model with unknown model parameters. Furthermore, in a more challenging application that we call meta-modulation, which is a more complex blind deconvolution problem with sophisticated system evolution equations, the proposed method performs satisfactorily and achieves an exciting result for high efficiency communication.

Keywords: sequential learning, sequential Monte Carlo, Kullback-Leibler divergence, blind deconvolution, meta-modulation

1. Introduction

State-space model, a class of probabilistic graphical model (Koller and Friedman, 2009) that describes the dependence between the unobserved state variable and the observed measurement, is a fundamental model for statistical inference with diversely applications in fields like statistics, econometrics, and information engineering (West and Harrison, 1997; Cappé et al., 2005). A linear or 'weakly' nonlinear state space model with Gaussian noise can be easily solved using Kalman filter and its derivatives. In nonlinear non-Gaussian scenarios, particle methods (also known as Sequential Monte Carlo methods) have been proven to be the most successful approach for the numerical approximation and Bayesian inference of the unknown state, because of their simplicity, flexibility, and ease of implementation (Gordon et al., 1993; Liu and Chen, 1998; Doucet et al., 2000; Liu, 2001; Doucet et al., 2001). In most of previous research, the inference has been focused on filtering or smoothing for the state, with the assumption that the model parameter in the state-space model are known. However, in practical situations, the model may contain unknown parameter that need to be estimated simultaneously with the state, and it might even be the case that the inference of the model parameter is the primary problem of interest.

One early and straightforward way to deal with this problem is to extend the original state to an augmented state that includes the state and the parameter together, and then to apply a standard particle filter to perform inference for both of them. However, this naive approach has been recognized not to be efficient, because the parameter space is not well explored (Kitagawa, 1998; Liu and West, 2001). Consequently, various improve methods has been developed over the past fifteen years (Refer to Kantas et al. (2015) for thoroughly review): maximum likelihood methods have been developed with different Monte Carlo evaluations of the likelihood of the model parameter (Hürzeler and Künsch, 2001; DeJong et al. 2013), and gradient based optimization (Ionides et al., 2006; Ionides et al., 2011) or expectation maximization methods (Andrieu et al., 2005; Cappé, 2009) have been introduced for an on-line or off-line estimation of the model parameter. The maximum likelihood approach generally converges rather slowly, but it may be a good choice for large data sets because of its low complexity; Bayesian methods apply directly to the augmented states and Markov chain Monte Carlo steps are utilized to improve the inference/estimation of the model parameter (Gilks and Berzuini, 2001; Fearnhead, 2002; Andrieu et al., 2010). In the inspiring work by(Liu and West, 2001), an artificial dynamic is introduced to the static model parameter, and a kernel density estimation method is proposed to capture the density of the parameters. The significant idea of these authors is to use a shrinkage strategy for kernel locations and variances inflation, which therefore removes the problem of information loss over time (West, 1993). However, the learning method proposed in this work seems to be ad hoc and weak in explaining the underlying driving strategy. In this paper, a stochastic algorithm is proposed that

tunes a flexible proposal mixture distribution to minimize the Kullback-Leibler divergence between the distribution of the unknown model parameter and the distribution of the proposed mixture: this leads to an online learning method for the numerical approximation of the distribution of the model parameter. We derive the sequential learning method by using a truncated Dirichlet processes normal mixture, and present the general algorithm using a proposed learning approach under a framework of the auxiliary particle filtering. The proposed algorithm is verified and compared with other related algorithms in a blind deconvolution problem, which is a typical state-space model with unknown model parameter (Liu and Chen, 1995). Furthermore, a more novel and challenging application is presented, which is essentially a more complex blind deconvolution problem with sophisticated system evolution equation. The proposed method performs satisfactorily in dealing with the demodulation of this system and obtains an exciting result for high efficiency communication.

The paper is organized as follows. In Section 2 we present the general state space model and the auxiliary particle filter algorithm for the Bayesian inference of this model. In Section 3 we introduce the sequential learning method and a general framework for joint parameter and state estimation. In Section 4 we verify the proposed method in a simple blind deconvolution and a more channelling problem in wireless communication. Finally, we summarize our methods in Section 5.

2. Problem Statement and Particle Filter Algorithm

A state-space model can be defined by the following two processes: the Markovian evolution equation, or state equation,

$$\mathbf{x}_t = f(\mathbf{x}_{t-1}, \theta) + w_t \tag{1}$$

defining the transition density $p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)$, in which the state vector at time t is $\mathbf{x}_t$, the model parameter vector is θ, $f(\cdot)$ is the system evolution function and w_t is the system noise; and the observation equation

$$\mathbf{y}_t = h(\mathbf{x}_t, \theta) + v_t \tag{2}$$

defining the observation density $p(\mathbf{y}_t|\mathbf{x}_{t-1}, \theta)$, where $h(\cdot)$ is the observation function and v_t is the observation noise. The above expressions covers a very broad class of interesting dynamic models (West and Harrison 1997).

With a known parameter θ the task of sequential Bayesian inference is to estimate the posterior distribution $p(\mathbf{x}_t|\mathbf{y}_{0:t}, \theta)$. However, to apply particle methods, a more general approach is to estimate the sequence of joint posteriors $p(\mathbf{x}_{0:t}|\mathbf{y}_{0:t}, \theta)$ recursively, which leads to the following fundamental recursions: for $t \geq 1$,

$$
\begin{align}
p(\mathbf{x}_{0:t}|\mathbf{y}_{0:t}, \theta) &= \frac{p(\mathbf{x}_{0:t}, \mathbf{y}_{0:t}|\theta)}{p(\mathbf{y}_{0:t}|\theta)} \tag{3} \\
&= \frac{p(\mathbf{x}_t, \mathbf{y}_t|\mathbf{x}_{0:t-1}, \mathbf{y}_{0:t-1}, \theta)p(\mathbf{x}_{0:t-1}|\mathbf{y}_{0:t-1}, \theta)p(\mathbf{y}_{0:t-1}|\theta)}{p(\mathbf{y}_t|\mathbf{y}_{0:t-1}, \theta)p(\mathbf{y}_{0:t-1}|\theta)} \tag{4} \\
&= \frac{p(\mathbf{x}_t, \mathbf{y}_t|\mathbf{x}_{0:t-1}, \mathbf{y}_{0:t-1}, \theta)p(\mathbf{x}_{0:t-1}|\mathbf{y}_{0:t-1}, \theta)}{p(\mathbf{y}_t|\mathbf{y}_{0:t-1}, \theta)} \tag{5} \\
&= p(\mathbf{x}_{0:t-1}|\mathbf{y}_{0:t-1}, \theta)\frac{p(\mathbf{y}_t|\mathbf{x}_{0:t}, \mathbf{y}_{0:t-1}, \theta)p(\mathbf{x}_t|\mathbf{x}_{0:t-1}, \mathbf{y}_{0:t-1}, \theta)}{p(\mathbf{y}_t|\mathbf{y}_{0:t-1}, \theta)} \tag{6} \\
&= p(\mathbf{x}_{0:t-1}|\mathbf{y}_{0:t-1}, \theta)\frac{p(\mathbf{y}_t|\mathbf{x}_t, \theta)p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)}{p(\mathbf{y}_t|\mathbf{y}_{0:t-1}, \theta)} \tag{7}
\end{align}
$$

where $p(\mathbf{y}_t|\mathbf{y}_{0:t-1}, \theta) = \int p(\mathbf{x}_{t-1}|\mathbf{y}_{0:t-1}, \theta)p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)p(\mathbf{y}_t|\mathbf{x}_t, \theta)d\mathbf{x}_{t-1:t}$, $p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)$ represents the system dynamic in (1), and $p(\mathbf{y}_t|\mathbf{x}_t, \theta)$ refers to the likelihood function obtained in (2). The recursion is initialised with some distribution, for example, $p(\mathbf{x}_0)$.

Particle filtering is a class of importance sampling and resampling techniques designed to give a numerical approximation for the recursions in (7). We present the auxiliary particle filter (APF) (Pitt and Shephard, 1999) here, as it covers a class of particle filter algorithms (Kantas et al., 2015) and is widely used in parameter and state estimation (Liu and West, 2001; Flury and Shephard, 2011). Let the proposal be $q(\mathbf{x}_t, \mathbf{y}_t|\mathbf{x}_{t-1}, \phi) = q(\mathbf{x}_t|\mathbf{y}_t, \mathbf{x}_{t-1}, \phi)q(\mathbf{y}_t|\mathbf{x}_{t-1}, \phi)$, where $q(\mathbf{x}_t|\mathbf{y}_t, \mathbf{x}_{t-1}, \phi)$ is a probability density function which is easy to sample from and $q(\mathbf{y}_t|\mathbf{x}_{t-1}, \phi)$ is a nonnegative function that can be evaluated. The auxiliary particle filter sequentially draws samples from $q(\mathbf{x}_t|\mathbf{y}_t, \mathbf{x}_{t-1}, \phi)$ and calculates the following importance weights,

$$w_t(\mathbf{x}_{t-1:t}) = \frac{p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)p(\mathbf{y}_t|\mathbf{x}_t, \theta)}{q(\mathbf{x}_t, \mathbf{y}_t|\mathbf{x}_{t-1}, \phi)}. \tag{8}$$

For $t \geq 1$ the APF algorithm can be summarized as follows,

At iteration $t \geq 1$, for all $i = 1, ..., N$:

Step 1:   Sample $\mathbf{x}_t^{(i)} \sim q(\mathbf{x}_t | \mathbf{y}_t, \mathbf{x}_{t-1}^{(i)}, \phi)$ and set $\mathbf{x}_{0:t}^{(i)} = [\mathbf{x}_{0:t-1}^{(i)}, \mathbf{x}_t^{(i)}]$;

Step 2:   Compute the weights $\widetilde{w}_t^{(i)} \propto w_t(\mathbf{x}_{t-1:t}^{(i)}) q(\mathbf{y}_t | \mathbf{x}_{t-1}^{(i)}, \phi)$, where $\widetilde{w}_t^{(i)}$ is normalized as $\sum_{i=1}^{N} \widetilde{w}_t^{(i)} = 1$;

Step 3:   Resample $\mathbf{x}_{0:t}^{(i)} \sim \sum_{i=1}^{N} \widetilde{w}_t^{(i)} \delta(\mathbf{x}_{0:t}^{(i)})$.

In many applications, the model parameter θ is unknown and it might even be of interest for inference. A straightforward solution is to define an extended state that includes the state $\mathbf{x}_t$ and the parameter θ together. For example in (Liu and West, 2001), an 'artificial evolution' equation for the model parameter θ is introduced: $\theta_t = \theta_{t-1} + \varepsilon$, where $\varepsilon \sim N(0, \Sigma_t)$ with some specified variance matrix Σ_t. A kernel density estimation method is proposed to capture the density of the parameter θ_t. However, the method of the kernel density estimation is critical in such a problem. As studied by Liu and West (2001), they use a shrinkage strategy for kernel locations and variances estimation in order to removes the over-dispersion of the variances. Although this method works satisfactorily with a careful chosen shrinkage strategy, more efficient approach are still worth for further exploration (Kantas et al., 2015).

3. Sequential Learning and Joint Parameter and State Estimation

3.1 Sequential Learning

A sequential learning method is proposed here to deal with the online learning of the unknown distribution of the model parameter. Assume that the parameter vector θ is subject to some unknown distribution $\pi(\theta)$. We propose to utilize a parametric distribution $q(\theta; \psi)$ with the controlling parameter ψ to approximate the unknown $\pi(\theta)$. To measure the distance between these two distributions, the broadly used Kullback-Leibler divergence (KL-divergence) is employed here. The idea of the proposed sequential learning approach is to tune the parameter ψ by learning from samples of $\pi(\theta)$, and enable $q(\theta; \psi)$ to approximate $\pi(\theta)$ in the sense of minimizing the KL-divergence $\mathcal{D}[\pi(\cdot) \| q(\cdot)] = \mathbb{E}_{\pi}\left[\log \frac{\pi(\cdot)}{q(\cdot)}\right]$. Other criterion like moment matching can also be applied to measure the closeness of two distributions (Ji, 2006), but it is not straightforward to derive the algorithm to find the optimal controlling parameter of the proposal distribution. Under the measurement of KL-divergence, the optimal parameter ψ^* which minimizes $\mathcal{D}\left[\pi(\theta) \| q(\theta; \psi)\right]$ can be obtained by finding the root of the derivative of $\mathcal{D}$ (if exists):

$$h(\psi) = - \int \pi(\theta) \frac{\partial}{\partial \psi}\left(\log \frac{\pi(\theta)}{q(\theta; \psi)}\right) d\theta = 0. \tag{9}$$

The closed-form solution of the integral equation (9) is generally intractable, as $h(\psi)$ involves an integral with unknown distribution $\pi(\theta)$. However, suppose that we can obtain samples $\theta^{(1:N)}$ from $\pi(\theta)$, denote $f(\theta, \psi) = \frac{\partial}{\partial \psi}\left(\log \frac{\pi(\theta)}{q(\theta; \psi)}\right)$, then we can numerically evaluate $h(\psi)$ by Monte Carlo integration:

$$\hat{h}(\theta^{(1:N)}; \psi) = \frac{1}{N} \sum_{i=1}^{N} f(\theta^{(i)}, \psi) \tag{10}$$

where $\hat{h}(\theta^{(1:N)}; \psi)$ can be viewed as a noisy 'observation' of $h(\psi)$. One available approach for obtaining roots of $h(\psi) = 0$ when we only have noisy evaluations of $h(\psi)$ is the Stochastic Approximation (SA) algorithm (Kusher and Yin, 1997). The SA algorithm iteratively updates ψ to approximate its optimal values by the following formula:

$$\begin{aligned} \psi_t &= \psi_{t-1} + r_t \left[h(\psi_{t-1}) + \xi_t\right] \\ &= \psi_{t-1} + r_t \, \hat{h}(\theta_{t-1}^{(1:N)}; \psi_{t-1}) \end{aligned} \tag{11}$$

where t is an iterative index, $\{\xi_t\}$ is a sequence of 'noise' (thus the Monte Carlo estimation $\hat{h}(\theta_t^{(1:N)}; \psi_t)$ can be interpreted as the ground true $h(\psi_t)$ plus noise ξ_t), and $\{r_t\}$ is a sequence of decreasing step-sizes satisfying $\sum_t r_t = \infty$ and $\sum_t r_t^2 < \infty$. The proof of covergence of this Monte Carlo estimation based sequential learning algorithm is detailed in the appendix.

We assume that $q(\theta; \psi)$ is a truncated Dirichlet process (TDP) normal mixture (Ishwaran and James, 2001; Ji, 2009). Let ψ denotes the set of controlling parameters (V_k, μ_k, Σ_k), a TDP normal mixture can be expressed as, $q(\theta; \psi) = \sum_{k=1}^{K} \omega_k N(\cdot | \mu_k, \Sigma_k)$, where $\omega_k = V_k \prod_{l=1}^{k-1}(1 - V_l)$ and $V_l \in (0, 1)$ which can be initially drawn from a beta distribution $Beta(\alpha_0, \beta_0)$. The the partial derivative of $\mathcal{D}[\pi(\theta) \| q(\theta; \psi)]$ with respect to V_k, μ_k and Σ_k can be derived respectively as follows (refer to Ji (2009) for more details of derivations),

$$h_{V_k}(\theta; \psi) = \int \pi(\theta) \frac{- \sum_{l=k+1}^{K} V_l \prod_{\iota \le l-1, \iota \ne k}(1 - V_\iota) q(\theta | \mu_l, \Sigma_l) + \prod_{l=1}^{k-1}(1 - V_l) q(\theta | \mu_k, \Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta | \mu_m, \Sigma_m)} d\theta \tag{12}$$

$$H_{\mu_k}(\theta;\psi) \propto \int \pi(\theta) \frac{\omega_k q(\theta;\mu_k,\Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta;\mu_m,\Sigma_m)} \times (\theta - \mu_k)\, d\theta \tag{13}$$

$$H_{\Sigma_k}(\theta;\psi) \propto \int \pi(\theta) \frac{\omega_k q(\theta;\mu_k,\Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta;\mu_m,\Sigma_m)} \times \left((\theta - \mu_k)(\theta - \mu_k)^T - \Sigma_k \right) d\theta \tag{14}$$

Given the samples $\Theta_t = \{\theta_t^{(i)}\}_{i=1}^{N}$ from $\pi(\theta)$, the Monte Carlo approximation of these partial derivatives is

$$H_{V_k}(\Theta_t;\psi) = \frac{1}{N} \sum_{i=1}^{N} \frac{-\sum_{l=k+1}^{K} V_l \prod_{\iota \le l-1, \iota \ne k} (1 - V_\iota) q(\theta_t^{(i)}|\mu_l,\Sigma_l) + \prod_{l=1}^{k-1} (1 - V_l) q(\theta_t^{(i)}|\mu_k,\Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta_t^{(i)}|\mu_m,\Sigma_m)} \tag{15}$$

$$H_{\mu_k}(\Theta_t;\psi) \propto \frac{1}{N} \sum_{i=1}^{N} \frac{\omega_k q(\theta_t^{(i)};\mu_k,\Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta_t^{(i)};\mu_m,\Sigma_m)} \times \left(\theta_t^{(i)} - \mu_k \right) \tag{16}$$

$$H_{\Sigma_k}(\Theta_t;\psi) \propto \frac{1}{N} \sum_{i=1}^{N} \frac{\omega_k q(\theta_t^{(i)};\mu_k,\Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta_t^{(i)};\mu_m,\Sigma_m)} \times \left(\left(\theta_t^{(i)} - \mu_k \right)\left(\theta_t^{(i)} - \mu_k \right)^T - \Sigma_k \right) \tag{17}$$

3.2 Joint Parameter and State Estimation

Let us return to the inference of the augmented states under a particle filter framework. Assuming the particle samples $(\mathbf{x}_t^{(i)}, \theta_t^{(i)})$ and weights $w_t^{(i)}$ ($i = 1, ..., N$) are available to represent the joint posterior $p(\mathbf{x}_t, \theta|\mathbf{y}_t)$, with the sequential learning algorithm to approach the unknown distribution of θ by samples $\theta_t^{(i)}$ ($i = 1, ..., N$), we now have an extended version of the auxiliary particle filter algorithm, incorporating the parameter and the state estimation together. The resulting general algorithm is presented as follows:

Step 1: At iteration t, the sampled N particles containt $\left\{ \mathbf{x}_t^{(i)}, \theta_t^{(i)}, c_t^{(i)} \right\}_{i=1}^{N}$ (where $c_t^{(i)}$ is the label of the TDP normal mixture component which $\theta_t^{(i)}$ belongs to), with its weights $w_t^{(i)}$. Calculate the auxiliary variable

$$\phi_{t+1}^{(i)} = E(\mathbf{x}_{t+1}|\mathbf{x}_t^{(i)}, \theta_t^{(i)}).$$

Step 2: Sample the auxiliary index from the set $\{1, ..., N\}$ with probabilities proportion to

$$g_{t+1}^{(i)} \propto w_t^{(i)} p(\mathbf{y}_{t+1}|\phi_{t+1}^{(i)}, \theta_t^{(i)});$$

obtain new index set j, and new label set $c_t^{(j)}$.

Step 3: Sample a new parameter vector $\theta_{t+1}^{(j)}$ from the $k = c_t^{(j)}$ normal components of the mixture model

$$\theta_{t+1}^{(j)} \sim N(\cdot|\mu_{k,t}, \Sigma_{k,t}).$$

Step 4: Sample the new state vector $\mathbf{x}_{t+1}^{(j)}$, from the system equation

$$\mathbf{x}_{t+1}^{(j)} \sim p(\cdot|\mathbf{x}_t^{(j)}, \theta_{t+1}^{(j)}).$$

Step 5: Calculate the importance weights

$$w_{t+1}^{(j)} \propto \frac{p(\mathbf{y}_{t+1}|\mathbf{x}_{t+1}^{(j)}, \theta_{t+1}^{(j)})}{p(\mathbf{y}_{t+1}|\phi_{t+1}^{(j)}, \theta_t^{(j)})}$$

Step 6:   Update the parameters of the TDP mixture V_k, ω_k, μ_k and Σ_k (for $k = 1, ..., K$) using $\left\{w_{t+1}^{(j)}, \theta_{t+1}^{(j)}\right\}_{j=1}^{N}$ as follows,

$$V_{k,t+1} \quad = \quad V_{k,t} + r_{t+1} H_{V_k}(\Theta_t; \psi_t) \tag{18a}$$

$$\omega_{k,t+1} \quad = \quad V_{k,t+1} \prod_{l=1}^{k-1}(1 - V_{l,t+1}) \tag{18b}$$

$$\mu_{k,t+1} \quad = \quad \mu_{k,t} + r'_{t+1} \sum_{j=1}^{N} w_{t+1}^{(j)} \alpha_{k,t+1}^{(j)} \left(\theta_{t+1}^{(j)} - \mu_{k,t}\right) \tag{18c}$$

$$\Sigma_{k,t+1} \quad = \quad \Sigma_{k,t} + r'_{t+1} \sum_{j=1}^{N} w_{t+1}^{(j)} \alpha_{k,t+1}^{(j)} \left((\theta_{t+1}^{(j)} - \mu_{k,t})(\theta_{t+1}^{(j)} - \mu_{k,t})^T - \Sigma_{k,t}\right) \tag{18d}$$

where $H_{V_k}(\Theta_t; \psi_t) = \sum_{i=1}^{N} w_{t+1}^{(j)} \dfrac{-\sum_{l=k+1}^{K} V_l \prod_{l \leq l-1, l \neq k}(1-V_l)q(\theta_t^{(j)}|\mu_l, \Sigma_l) + \prod_{l=1}^{k-1}(1-V_l)q(\theta_t^{(j)}|\mu_k, \Sigma_k)}{\sum_{m=1}^{K} \omega_m q(\theta_t^{(j)}|\mu_m, \Sigma_m)}$, $\alpha_{k,t+1}^{(j)} = \dfrac{\omega_{k,t} q(\theta_{t+1}^{(j)}; \mu_{k,t}, \Sigma_{k,t})}{\sum_{m=1}^{K} \omega_{m,t} q(\theta_{t+1}^{(j)}; \mu_{m,t}, \Sigma_{m,t})}$, r_{t+1} and r'_{t+1} are the step-sizes in the stochastic approximation algorithm.

Step 7:   If the stopping criterion is satisfied, then stop; otherwise, set $t := t + 1$ and go back to step 1.

3. Simulation Study

3.1 Blind Deconvolution

Blind deconvolution of source signals is a subject that has been widely studied and applied in various fields, such as wireless communications, sonar and radar systems, audio and acoustics, and image processing (Benveniste and Goursat, 1984; Haykin, 1994). Take the wireless communication over a multipath fading channel as an example (Chen and Liu, 2000; Ali et al., 2002), the unknown information signal propagates through the channel which mixes the signal as it passes through a filter. The blind deconvolution problem is to recover the independent sources from the mixed signal without any priori knowledge of the original signal and the conditions of the channel. This problem is a typical state-space model problem, and can be expressed as follows,

$$x_t \quad = \quad \sum_{i=0}^{q} h_i s_{t-i} \tag{19}$$

$$y_t \quad = \quad x_t + \epsilon_t \tag{20}$$

where the information bits $s_t \in \{-1, 1\}$ are independent, identically distributed (i.i.d) samples from a *Bernoulli*(0.5), $[h0, h_1, ..., h_q]^T$ is a set of coefficients representing the channel conditions, q is the number of the multipath, y_t is the received signal for decoding and the noise ϵ are i.i.d. normal with mean zero and variance σ.

In the simulation study, we consider the following settings: the number of particles in the particle filter is 200, the number of components in the TDP normal mixture is 20, the initial values of μ_k are randomly chosen from $(-0.5, 1)$ and Σ_k are $0.1 * \mathbf{I}$ where $\mathbf{I}$ is an identity matrix. The noise variance σ is an important parameter in the blind deconvolution problem, where we use the signal to noise ratios(SNR) to measure the noise level. For the first study, the SNR is 14 dB and the coefficients of the channel are shown in Table 1. We run 10 trials, and each trial has 10, 000 information bits. The sequential learning algorithm outputs an online estimation of the coefficients $\theta \triangleq [h_0, h_1, h_2]^T$. At each iteration we estimate the mean value of θ using the particles $\left\{w_t^{(j)}, \theta_t^{(j)}\right\}_{j=1}^{200}$. The resulting trajectories from 1 to 2000 are shown in Figure 1, which demonstrates the convergence rate and stability of the sequential learning algorithm. The mean values of each h_i over a long period (iterations 1, 000 to 10, 000) are shown in Table 1, which show the high accuracy of our algorithm in estimating the model parameter. The performance of blind deconvlution in the above communication system can be well measured using the bit error rate (BER) of different SNRs (Chen and Liu, 2000; Ali et al., 2002). We compare our proposed method (SL+APF) with a naive bootstrap particle filter (BPF) with augmented state (x, θ), a BPF with known θ, an APF with augmented state (x, θ), an APF with known θ and a matched filter bound (MFB) which represents a reference line for the BER in Gaussian noise channel (Ali et al., 2002). The MFB can be calculated as follows (Ali et al., 2002): $1 - \int_{-\infty}^{p} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx$, where $p = \frac{\sqrt{\sum_{i=0}^{q} h_i^2}}{\sigma}$. Figure 2 shows the BERs for different SNRs (in dB). The APF based methods can generally outperform the BPF based methods. The proposed SL+APF method outperforms the APF with augmented states, and the resulting BER is close to the one evaluated by the APF when θ is known. Moreover, APF with augmented states has a apparent higher SNR requirement for low BERs below 10^{-3}, while our proposed method overcomes this drawback because of the efficient learning for channel coefficients.

Table 1. The coefficients of channel response and the estimation via sequential learning algorithm

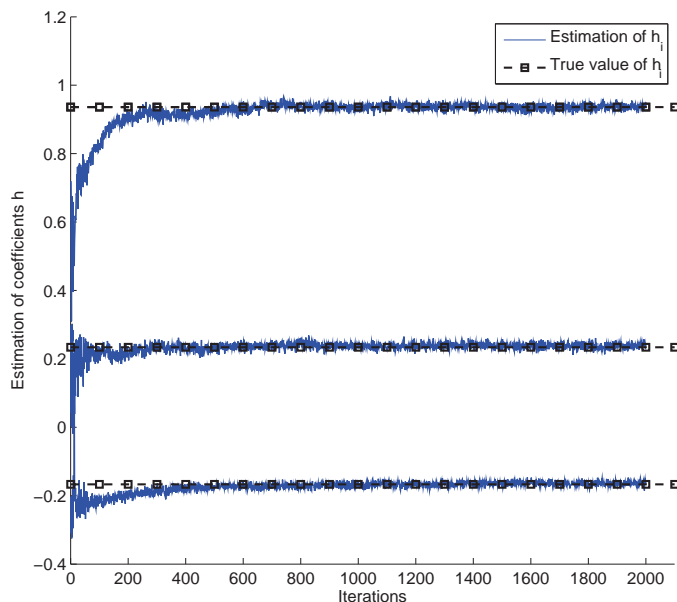| Coefficients	h_0	h_1	h_2
True value	0.9460	0.2340	-0.1670
Estimation	0.9450	0.2322	-0.1669

Figure 1. Trajectory plots of the mean value of $\theta = [h_0, h_1, h_2]^T$ estimated using the particles $\left\{w_t^{(j)}, \theta_t^{(j)}\right\}_{j=1}^{200}$ from iteration 0 to 2000. The estimation of each h_i comes close to the true value with in a few hundreds iterations and becomes stable since then.

3.2 Meta-modulation

In this example we consider a realistic state space model with sophisticated system evolution equations. In the above example, only the channel response is considered as the mixing mechanism, which is expressed as the coefficients **h**. Here, we consider a more general situation, where the impulse response consists of all the filters and the channel response during the entire transmission. Let $\mathbf{h}_t = [h_0, ..., h_{q_t}]^T$ represents all the filter responses in the transmitter, such as the pulse shaping filter, the transmitter filter and etc. Let $\mathbf{h}_c = [h_0, ..., h_{q_c}]^T$ be the channel response, and let $\mathbf{h}_r = [h_0, ..., h_{q_r}]^T$ be the response of all the filters in the receiver, such as the receiver filter, the matched filter and etc.

$$x_t = \sum_{q=0}^{Q} h_q s_{t-q} \tag{21}$$

$$y_t = x_t + \epsilon_t \tag{22}$$

where $\mathbf{h} = [h_0, h_2, ..., h_Q]^T$ represents the convolution of all the impulse responses $\mathbf{h} = \mathbf{h}_t \otimes \mathbf{h}_c \otimes \mathbf{h}_r$, and $Q = q_t + q_c + q_r$ is the total number of the filter coefficients.

Let $S = [s_0, s_\tau, s_{2\tau}, ...]^T$ denotes the transmitted i.i.d. bit sequence. In traditional communication system, S will be N times upsampled to $S' = [s_0, 0, ..., 0, s_\tau, 0, ..., 0, s_{2\tau}...]^T$, and then S' passes through a pulse shaping filter $h(n\tau/N)$ ($n = 0, ..., N - 1$) with duration τ. The resulting convolution signal will be

$$\mathbf{x} = [s_0 h((N - 1)\tau/N), s_0 h((N - 2)\tau/N), ..., s_0 h(0), s_\tau h((N - 1)\tau/N), s_\tau h((N - 2)\tau/N), ..., s_\tau h(0), ...]^T. \tag{23}$$

To demodulate such a signal, one can apply the matched filter to the received noisy signal **y**, where $\mathbf{y} = \mathbf{x} + \epsilon$, and then decide the value of each s_i by using the signal detection theory. Obviously, such a demodulation is not a blind deconvolution problem. On the other hand, consider the scenario that S is transmitted Q times faster, then S becomes
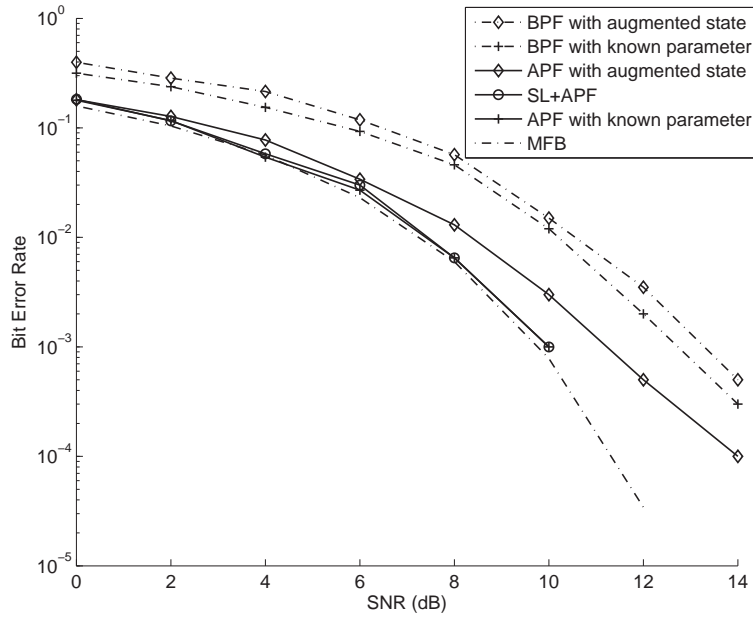
Figure 2. Comparisons of BERs for different methods. Both APF with known parameter and our proposed SL+APF algorithms deconvolute all the 10,000 bits correctly when SNR ≥ 12dB, so no points appear in 12dB and 14dB for these two curves.

$S' = [s_0, s_{\tau/Q}, ..., s_{(Q-1)\tau/Q}, s_\tau, ...]^T$. Let S' passes through a pulse shaping filter $h(q\tau/Q)$ $(q = 0, ..., Q-1)$ with duration τ, then we get a truly convoluted signal

$$\tilde{\mathbf{x}} = [s_0 h((Q-1)\tau/Q), s_0 h((Q-2)\tau/Q) + s_{\tau/Q} h((Q-1)\tau/Q), ..., \sum_{q=0}^{Q} s_{(n-q)\tau/Q} h(q\tau/Q), ...]^T. \qquad (24)$$

Consequently this communication system violates the Nyquist intersymbol interference criterion, and traditional matched filter and detection approaches fail to demodulated such signal, however, we can appeal to the blind deconvolution techniques to separate the information bits at the receiver. The notable advantage of this system is that the bandwidth of modulated signal $\tilde{\mathbf{x}}$, which depends mainly on the pulse shaping filter $h(q\tau/Q)$, is comparable with the bandwidth of $\mathbf{x}$ as the pulse duration τ is the same, but the information sequence S has been managed to transmitted Q times faster.

Furthermore, it can be verified that this system essentially belongs to a parallel Gaussian channel model: let the vector $S_{Q\times 1}$ represents Q bits information, $H_{(2Q-1)\times Q}$ be a matrix that each row is the filter coefficients $\mathbf{h}^T$ shifted with one time slot delay, then the received signal $Y_{(2Q-1)\times 1}$ is given by the matrix format $Y = HS + \epsilon$. It is trivial to find that the rank of matrix $H_{(2Q-1)\times Q}$ is Q. Apply a singular value decomposition, $H = U\Lambda V$, where both $U_{(2Q-1)\times(2Q-1)}$ and $V_{Q\times Q}$ are the unitary matrix and $\Lambda_{(2Q-1)\times Q}$ is a diagonal matrix with rank Q, then

$$Y' = \Lambda S' + \epsilon' \qquad (25)$$

where $Y' = U^{-1}Y$, $S' = VS$ and $\epsilon' = U^{-1}\epsilon$. According to (25) , this system is equivalent to Q independent parallel channels (Tse and Viswanath 2005; Cover and Thomas, 2006). Consequentially, the transmission rate R of this communication system will be bounded by the capacity of a parallel Gaussian channel model: $R \leq C = \frac{Q}{2} \log_2(1 + \text{SNR}/Q)$ bits per transmission. This high efficient communication strategy is proposed to be named as 'meta-modulation': it utilizes ordinary modulation framework but allows heavy interaction such as convolution in certain domain of the information bits, and therefore format an independent parallel channel to achieve an extraordinary high spectrum efficiency.

In this simulation study, we consider the case that Q is sufficient large, then it is not trivial to demodulate the convolutional signal of (21) particulary some of the coefficients are unknown. For example, the Viterbi algorithm can be modified to deal with deconvolution or blind deconvoultion (Ali et al., 2001; Li, 2014), but it becomes infeasible for large Q as the potential paths in Viterbi algorithm increase exponentially. However, our proposed SL+APF method seems well suite for the deconvolution of this meta-modulation system. The simulation is set as follows, h_t is a Gaussian filter with $q_t = 100$, h_r is also a Gaussian filter with $q_r = 100$, $h_c = [0.925, 0.430, -0.201, -0.117, 0.091]$. Parameter settings of the algorithm are
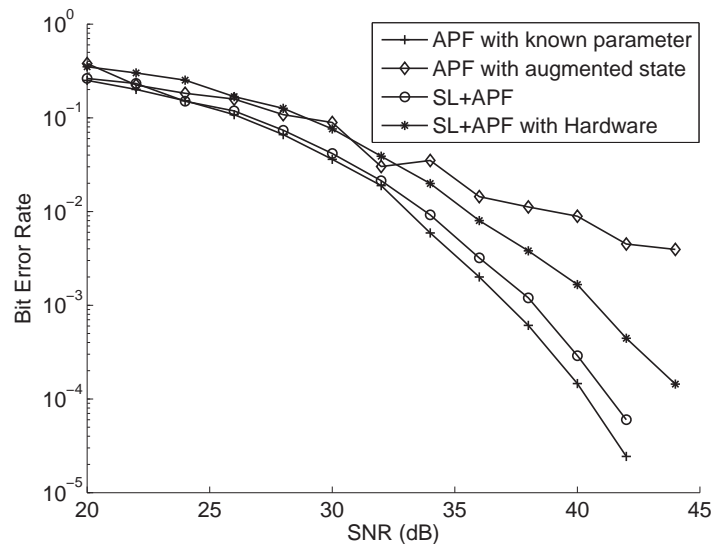
Figure 3. Comparisons of BERs for different methods. Both APF with known parameter and our proposed SL+APF algorithms deconvolute all the 10^5 bits correctly when SNR > 42dB, so no points appear in 44 dB for these two curves.

kept the same with example 3.1 except we simulated 10^5 bits. Coefficients of h_t and h_r are predefined and known, while the channel response h_c are unknown at the receiver. We compare our proposed method with an APF with augmented state and an APF with known channel coefficients. The BER plots in Figure 3 shows that the proposed SL+APF algorithm still performs well, but this time the APF with augmented state method performs very unstable because of the insufficiency in exploration of the unknown channel coefficients. Simulation study also demonstrates that the proposed algorithm achieves the BER < 10^{-5} at SNR=44 dB, which is approximately 4 dB higher than the theoretical required SNR 40.1dB to reach the capacity 200 bits per transmission. We also verify that the meta-modulation works in real-world physical channels, we implemented the communication process on a standard verification system, the universal software radio peripheral (USRP). The received signal is demodulated by the SL+APF algorithm. As shown in Figure 3, the BER becomes slightly higher because of the quantization error in the real communication, however, hardware verification shows the potential in real-world applications. In conclusion, this is an exciting result, because not only such a high spectrum efficiency system has rarely been reported, but also the required SNR for demodulation is moderately low that can be satisfied in practice.

Acknowledgement

References

Ali, R. A., Murua, A., Richardson, T., Roy, S. (2002). Evaluation of Sequential Importance Sampling for blind deconvolution via a simulation study. *In Proc. 11th European Signal Processing Conference*, 1-4.

Andrieu, C., Doucet, A., & Holenstein, R. (2010) . Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, *72*, 269-342. http://dx.doi.org/10.1111/j.1467-9868.2009.00736.x

Andrieu, C., Doucet, A., & Tadić, V. B. (2005). Online parameter estimation in general state-space models. *In Proc. 44th IEEE Conf. on Decision and Control*, 332-337. http://dx.doi.org/10.1109/CDC.2005.1582177

Benveniste, A., & Goursat, M. (1984). Blind equalizers. *IEEE Transactions on Communications*, *32*, 871-883. http://dx.doi.org/10.1109/TCOM.1984.1096163

Cappé, O. (2009). Online sequential Monte Carlo EM algorithm. *In Proc. 15th IEEE Workshop on Statistical Signal Processing.* 37-40. http://dx.doi.org/10.1109/ssp.2009.5278646

Cappé, O., Moulines, E., & Rydén, T. (2005). *Inference in Hidden Markov Models.* Springer, New York.

Chen, R., & Liu, J. S. (2000). Mixture Kalman Filters. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, *62*, 493-508. http://dx.doi.org/10.1111/1467-9868.00246

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory*. John Wiley & Sons, Inc.

DeJong, D. N., Liesenfeld, R., Moura, G. V., Richard, J.-F., & Dharmarajan, H. (2013). Efficient likelihood evaluation of state-space representations. *Review of Economic Studies*, *80*, 538-567. http://dx.doi.org/10.1093/restud/rds040

Doucet, A., De Freitas, J. F. G., & Gordon, N. J., eds. (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York.

Doucet, A., Godsill, S. J., & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, *10*, 197-208. http://dx.doi.org/10.1023/A:1008935410038

Fearnhead, P. (2002). Markov chain Monte Carlo, sufficient statistics, and particle filters. *Journal of Computational and Graphical Statistics*, *11*, 848-862. http://dx.doi.org/10.1198/106186002835

Flury, T., & Shephard , N. (2011). Bayesian inference based only on simulated likelihood: Particle filter analysis of dynamic economic models. *Econometric Theory*, *27*, 933-956. http://dx.doi.org/10.1017/S0266466610000599

Gilks, W. R., & Berzuini, C. (2001). Following a moving target‐Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, *63*, 127-146. http://dx.doi.org/10.1111/1467-9868.00280

Gordon, N. J., Salmond, D. J., & Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F, Radar and Signal Processing*, *140*, 107-113. http://dx.doi.org/10.1049/ip-f-2.1993.0015

Haykin, S. (1994). *Blind Deconvolution*. Prentice-Hall, New Jersey.

Hürzeler, M., & Künsch , H. R. (2001). Approximating and maximising the likelihood for a general state-space model. *In Sequential Monte Carlo Methods in Practice. Stat. Eng. Inf. Sci*, 159-175. http://dx.doi.org/10.1007/978-1-4757-3437-9_8

Kantas, N., Doucet, A., Singh S. S., Maciejowski, J., & Chopin, N. (2015). On Particle Methods for Parameter Estimation in State-Space Models. *Statistical Science*, *30*(3), 328-351. http://dx.doi.org/10.1214/14-STS511

Kitagawa, G. (1998). A self-organizing state-space model.*Journal of the American Statistical Association*, *93*, 1203-1215. http://dx.doi.org/10.2307/2669862

Koller, D., & Friedman N. (2009) *Probabilistic Graphical Models*. Cambridge, MA: MIT Press.

Kushner, H. J., & Yin, G. G. (1997), *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, New York.

Ionides, E. L., Bretó, C., & King, A. A. (2006). Inference for nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA*, *103*, 18438-18443. http://dx.doi.org/10.1073/pnas.0603181103

Ishwaran, H., & James, L. (2001). Gibbs sampling methods for stick-breaking priors, *Journal of the American Statistical Association*, *96*, 161-173. http://dx.doi.org/10.1198/016214501750332758

Ji, C. (2006). *Adaptive Monte Carlo Methods for Bayesian Inference*, MPhil dissertation, University of Cambridge.

Ji, C. (2009). *Advances in Bayesian Modelling and Computation: Spatio-temporal Processes, Model Assessment and Adaptive MCMC*, PhD dissertation, Duke University.

Liu, J., & West, M. (2001). Combined parameter and state estimation in simulation-based filtering. *In Sequential Monte Carlo Methods in Practice*. Springer, New York. http://dx.doi.org/10.1007/978-1-4757-3437-9_10

Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer, New York.

Liu, J. S., & Chen, R. (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association*, *90*, 567-576. http://dx.doi.org/10.1080/01621459.1995.10476549

Liu, J. S., & Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, *93*, 1032-1044. http://dx.doi.org/10.1080/01621459.1998.10473765

Olsson, J., Cappé, O., Douc, R., & Moulines, E. (2008). Sequential Monte Carlo smoothing with application to parameter estimation in nonlinear state space models. *Bernoulli*, *14*, 155-179. http://dx.doi.org/10.3150/07-BEJ6150

Pitt, M., & Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters, *Journal of the American Statistical Association*, *94*, 590-599. http://dx.doi.org/10.1080/01621459.1999.10474153

Poyiadjis, G., Doucet, A., & Singh, S. S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, *98*, 65-80.

http://dx.doi.org/10.1093/biomet/asq062

Robert, C., & Casella, G. (1999). *Monte Carlo Statistical Methods*, 2nd ed. Springer, New York.

Wang, I. , Chong, E., and Kulkarni, S. (1995). On Equivalence of Some Noise Conditions for Stochastic Approximation Algorithms, *in Proceedings of the 34th Conference on Decision and Control.*

West, M. (1993). Approximating posterior distributions by mixtures, *Journal of Royal Statistical Society 55*, 409-422.

West, M., & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*, 2nd ed. Springer, New York.

Appendix: Proofs of Convergence of Monte Carlo SA Algorithm

Stochastic Approximation (SA) is a class of algorithms to finding the roots of possibly non-linear equation $h(\psi) = 0$, in the situation where only noisy measurements of $h(\psi)$ are available. In its simplest form, the Robbins-Monro algorithm is a recursive process as follows,

$$\psi^{(t+1)} = \psi^{(t)} + r^{(t+1)}\zeta^{(t+1)} \tag{26}$$

where $\left\{r^{(t)}, t \geq 1\right\}$ is a sequence of stepsizes which satisfies standard conditions: $\sum_{t=1}^{\infty} r^{(t)} = \infty$ and $\sum_{t=1}^{\infty} \left[r^{(t)}\right]^2 < \infty$ and for any $t \geq 1$, ζ is a noisy measurement of $h(\psi)$:

$$\zeta^{(t+1)} = h(\psi) + \xi^{(t+1)} \tag{27}$$

where $\left\{\xi^{(t)}, t \geq 1\right\}$ is the so called noise sequence.

In our case, we denote $f(\theta; \psi) = \frac{\partial}{\partial \psi}\left(\log \frac{\pi(\theta)}{q(\theta;\psi)}\right)$. Assume we have Monte Carlo samples $\{\theta^{(i)} : i = 1, \ldots, N\}$ from the distribution $\pi(\theta)$, then $h(\psi)$ can be evaluated by its Monte Carlo estimate,

$$\zeta(\psi) = \frac{1}{N}\sum_{i=1}^{N} f(\theta^{(i)}; \psi). \tag{28}$$

The *Central Limit Theorem* gives

$$\xi = [\zeta(\psi) - h(\psi)] \rightarrow Norm\left(0, \frac{\sigma^2}{N}\right), \quad \text{as} \quad N \rightarrow \infty \tag{29}$$

which implies that ξ is Gaussian noise, with mean zero and variance $\frac{\sigma^2}{N}$ where $\sigma^2 = \frac{1}{N-1}\sum_{i=1}^{N}(h(\psi) - \zeta(\psi))$ (Robert and Casella, 1999).

By using the iterative stochastic approximation method, we can estimate ψ iteratively by

$$\psi^{(t+1)} = \psi^{(t)} + r^{(t+1)}\zeta\left(\psi^{(t)}\right). \tag{30}$$

Assume that ψ_L is the root of equation $h(\psi) = 0$. We present a theorem to show that $\psi^{(t)} \rightarrow \psi_L$ in probability one as $t \rightarrow \infty$.

Theorem 1. *Consider the following conditions:*

(A1) *By the central limit theorem, $\xi(\psi) \rightarrow N\left(0, \frac{\sigma^2}{N}\right)$ in distribution and*

$$\int_{\Theta} h(\theta; \psi)^2 q(\theta; \psi) d\theta < \infty.$$

(A2) Ψ *is an open subset of R^{n_ψ}. The mean field $f : \Psi \rightarrow R^{n_\psi}$ is continuous and there exists a continuously differentiable function $w : \Psi \rightarrow [0, \infty)$ (with the convention $w(\psi) = \infty$ when $\psi \notin \Psi$) such that:*

1. *For any $M > 0$, the level set $W_M \equiv \{\psi \in \Psi, w(\psi) \leq M\} \subset \Psi$ is compact,*

2. *The set of stationary point(s) $\mathcal{L} \equiv \{\psi \in \Psi, \langle \nabla w(\psi), f(\psi) \rangle = 0\}$ belongs to the interior of Ψ,*

3. *For any $\psi \in \Psi$, $\langle \nabla w(\psi), f(\psi) \rangle \leq 0$ and the closure of $w(\mathcal{L})$ has an empty interior.*

(A3) *The sequence $\{r^{(t)}, t \geq 1\}$ is non-increasing, positive and*

$$\sum_{t=1}^{\infty} r^{(t)} = \infty \quad and \quad \sum_{t=1}^{\infty} \left[r^{(t)}\right]^2 < \infty. \tag{31}$$

Assume (A1-3). Then,

$$P\left[\lim_{t\to\infty} d\left(\psi^{(t)}, \mathcal{L}\right) = 0\right] = 1. \tag{32}$$

Proof: the recursion is expressed as follows

$$\psi^{(t+1)} = \psi^{(t)} + r^{(t+1)}\zeta^{(t+1)} = \psi^{(t)} + r^{(t+1)}f(\psi) + r^{(t+1)}\xi^{(t+1)}, \tag{33}$$

where $f(\psi)$ is the function of interest and $\xi^{(t+1)}$ is a random perturbation

$$\xi^{(t+1)} = f(\psi) - \zeta^{(t+1)} \tag{34}$$

$$= \int_{\Theta} h(\theta; \psi)d\theta - \frac{1}{N}\sum_{i=1}^{N} h(\theta^{(i)}; \psi). \tag{35}$$

Define $M_n = \sum_{t=1}^{n} r^{(t)}\xi^{(t)}$. Then

$$E[M_{n+1}|M_k, k \le n] = E\left[\sum_{t=1}^{n+1} r^{(t)}\xi^{(t)}|M_k, k \le n\right] \tag{36}$$

$$= E\left[\sum_{t=1}^{n} r^{(t)}\xi^{(t)}|M_k, k \le n\right] + E\left[r^{(t+1)}\xi^{(t+1)}|M_k, k \le n\right] \tag{37}$$

$$= M_n + r^{(t+1)}E\left[\xi^{(t+1)}\right]. \tag{38}$$

Since $\xi^{(t+1)} = \int_{\Theta} h(\theta; \psi)d\theta - \frac{1}{N}\sum_{i=1}^{N} h(\theta^{(i)}; \psi) \to 0$, almost sure (a.s.), as $N \to \infty$, $E[M_{n+1}|M_k, k \le n] \to M_n$ a.s. or with probability one. Therefore $\{M_n, n \ge 1\}$ is a F-martingale.

Then by the martingale inequality,

$$P\left\{\sup_{n\ge j\ge m} \left|M_j - M_m\right| \ge \mu\right\} \le \frac{E\left|\sum_{i=m}^{n-1} r^{(i)}\xi^{(i)}\right|}{\mu} \tag{39}$$

which implies

$$\lim_{m\to\infty} P\left\{\sup_{j\ge m} \left|M_j - M_m\right| \ge \mu\right\} = 0 \tag{40}$$

so that

$$\lim_{m\to\infty}\left(\sup_{m\le j\le m(n,T)} \left\|\sum_{i=n}^{j} r^{(i)}\xi^{(i)}\right\|\right) = 0 \quad \text{for all} \quad T > 0 \tag{41}$$

where $m(n, T) \equiv \max\left\{k : r^{(n)} + ... + r^{(k)} \le T\right\}$. This condition is called Kushner and Clark's condition, which is an important sufficient condition on the noise sequence for the convergence of stochastic approximation algorithms (Wang et al., 1995; Kushner and Yin, 1997). By the theorem in Chapter 5 of Kushner and Yin (1997), we can obtain

$$P\left[\lim_{t\to\infty} d\left(\psi^{(t)}, \mathcal{L}\right) = 0\right] = 1. \tag{42}$$

Copyrights