

Evaluation of Sequential Importance Sampling for Blind Deconvolution via a Simulation Study

R. A. Ali, A. Murua, T. Richardson
Department of Statistics, University of Washington
Seattle, WA, USA, 98195-4322

S. Roy
Department of Electrical Engineering, University of Washington
Seattle, WA, USA, 98195-2500
email: ayesha@stat.washington.edu

ABSTRACT

Sequential techniques for the canonical blind deconvolution problem have attracted the attention of computational Bayesians such as Liu and Chen (1995) who applied Sequential Importance Sampling (SIS) to this problem. Subsequently, several extensions have been proposed (e.g. Rejuvenation, Rejection Control, Fixed-Lag Smoothing, Metropolis-Hastings Importance Resampling, etc.) as improvements to SIS, but some of the drawbacks inherent in SIS persist. A comparison of variants of the Viterbi (VA), List Viterbi (LVA), BCJR (for Bahl, Cocke, Jelinek and Raviv) and SIS algorithms was conducted with inconclusive results. Although SIS can be helpful in certain circumstances, it shows signs of *instability*, and therefore, may not be useful in practice. In conclusion, one should be cautious in using SIS or Rejuvenation for blind deconvolution problems.

Keywords: blind deconvolution, sequential importance sampling, rejuvenation

1 Introduction

Consider the following linearly degraded system, which is often seen in signal processing:

$$y_t = \sum_{i=0}^q h_i x_{t-i} + \epsilon_t, \quad t = 1, 2, \dots, T \quad (1)$$

where $\epsilon_t \stackrel{i.i.d.}{\sim} N(0, \sigma^2)$, and the y_t are the observed output signals. The x_t are unobserved discrete input signals with known levels s_1, \dots, s_m , and the h_i are unknown coefficients for the the blurring mechanism. Our objective is to recover the input sequence x_t , as well as the coefficients h_0, \dots, h_q . We assume throughout that the system variance, σ^2 , is known since an approximate knowledge of σ^2 often suffices. A procedure that was sequential, adaptive and easy to implement would represent an ideal solution to this problem.

In the last few years several publications have appeared in both the Statistics and Signal Processing literature applying Bayesian techniques to blind deconvolution problems. However, from these papers, it is unclear what exactly is gained by these methods over what is

in use today. To gain some perspective on Sequential Importance Sampling based methods (Liu and Chen (1995)), we implemented naive versions of the Viterbi, List Viterbi, and BCJR algorithms with channel estimation (VCE, LVCE and BCJR H respectively) as a basis for comparison with the SIS-based methods. In particular, Recursive Least Squares (Seshadri and Sundberg (1994)) was used to estimate the channel and we implemented an estimation scheme that flipped between estimating the channel and estimating the input sequence. Details of these procedures are given in Ali et al. (2001). We believe such a study is valuable because (to the best of our knowledge) no such authoritative comparison has been previously reported.

2 Comparison of SIS to Traditional Methods for Blind Deconvolution

We consider the following blurring mechanism:

$$y_t = 0.407x_t + 0.815x_{t-1} + 0.407x_{t-2} + \epsilon_t.$$

where the output signals, x_t are i.i.d. *Bernoulli*($\frac{1}{2}$) from $\{-1, 1\}$, and the errors are i.i.d. *Normal* with mean zero and variance σ^2 . It is assumed that the data generating mechanism for the input signals, x_t , is known so no *Dirichlet* prior is incorporated into the decoding scheme for either SIS or Rejuvenation. Furthermore, it is assumed that the initial (x_{1-q}, \dots, x_0) bits are known. The various decoding methods mentioned in Section 1 were applied to this model for signal length $T = 1000$, and varying noise parameters σ^2 . The threshold used for Rejuvenation was 10% of the number of imputation sequences tracked. If $N = 1000$ imputed sequences were tracked, then the procedure rejuvenated whenever the Effective Sample Size (ESS) fell below 100.

Because each bit can take on one of two possible states, there is an obvious symmetry in the deconvolution problem. To circumvent this identifiability issue, it is assumed that the sign of the first coefficient is known. In particular, if the first coefficient of the channel estimate is negative, then the sign of every coefficient of the estimate is flipped, as is each bit of the corresponding estimate of the input sequence.

When the channel is symmetrical, as in this example, it is difficult to decode signals generated by that channel. Figure 1 shows the symbol or bit error rate (BER) curves, on a log-scale, for this channel using VA, LVA, VCE, LVCE, BCJR H, SIS, and Rejuvenation. This plot also shows the Matched Filter Bound (MFB) as a reference line for the BER. Recall that VA and LVA do not estimate the channel, but all the other methods do. In some sense, the gap between the VA line, and those corresponding to the blind decoding schemes represents the information lost by not knowing the channel. Each BER curve was averaged over 100 simulations.




Figure 1: Lag-2 Example with Channel (0.407, 0.815, 0.407), Signal Length $T = 1000$, $N = 200$ Imputation sequences tracked. SIS and Rejuv do not perform as well on this channel as on the original one. Rejuv again shows signs of instability. LVCE seems to perform rather well, and appears to be quite stable.

It is interesting to note that for this channel LVCE and VCE seem to outperform SIS, at least for SNR below 17db. Also, even though Rejuvenation seems unstable, the corresponding error rates do tend to be lower than that of the other methods.

2.1 Distribution of Error Rates

Figure 2 shows the distribution of the error rates for SIS and Rejuvenation for two channels. Each histogram is based on 100 simulations with signal length 1000 and tracking 200 imputation sequences. Each bar specifies the number of simulations, out of 100, that had the given percentage of errors. If we look at any histogram in Figure 2, we see that the number of simulations falling in the lower error rate intervals is usually higher for Rejuvenation than for Sequential Imputation. So, for the first channel with SNR equal to 1db (see Figure 2(e)), there were 22 simulations that had between 10 to 20% errors using SIS, and 73 simulations that had between




Figure 2: Histograms of Error Rates for SIS and Rejuvenation, $T=1000$, $N=200$. Note that using Rejuvenation results in more runs with fewer errors.

20 and 30% errors. The respective simulation frequencies were 76 and 16 for Rejuvenation. Resampling sequences, as prescribed by Rejuvenation, can help remove the imputed sequences resulting with high error rates. However, the earlier described instability in Rejuvenation persists. In Figure 2(e), there were 2 simulations under SIS that had more than 50% errors, and there were 6 such simulations under Rejuvenation. It is runs like these that make the BER curves look jumpy. Although Rejuvenation may help in improving the average BER, there is a tendency for more of these “outliers” to emerge. Histograms such as Figures 2(d) and (e) suggest that the average BER may sometimes be unrepresentative of performance on any sequence.

2.2 Distribution of Importance Weights

Ideally, any extension or variant of Rejuvenation would remove any instability from the procedure (evidenced by the lack of occasional runs with unexpectedly high

error rates). To study how these occasional “bad” runs (resulting with high error rates) differed from a typical “good” run (resulting with a low error rate), plots of ESS throughout different runs were examined.

In Figure 3 the ESS curve demonstrates some type of monotonic behavior for the “good” run. ESS starts at 200, the total number of sequences being tracked, and then decreases until it falls below the prescribed threshold (20), at which point the procedure rejuvenates and ESS is set back to 200. The corresponding plot for the “bad” run looks very different because the monotonic behavior between rejuvenation times is no longer apparent. This plot suggests that there is some extra variation in the distribution of the weights which thresholding ESS alone can not detect.

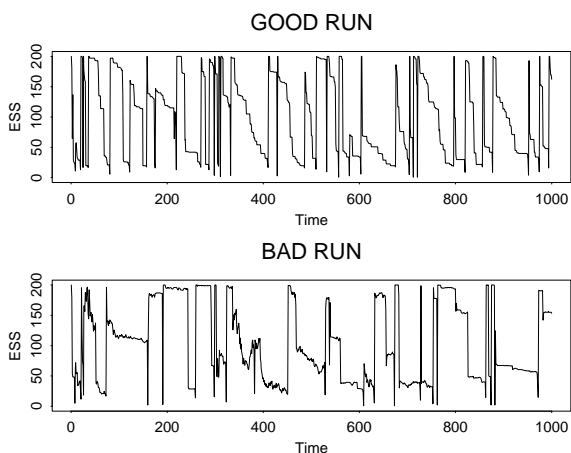


Figure 3: ESS Behavior through a “good” run (BER=0.0) and “bad” run (BER=0.46). Channel: (0.407, 0.815, 0.407), $T=1000$, $N=200$, SNR=13db. A “bad” run shows signs of non-monotonicity between rejuvenation times.

Figure 4 tracks the distribution of the weights throughout the runs depicted in Figure 3. Note that although the top line is the largest log weight associated with any of the 200 imputed sequences, it is not always the same sequence at all time points. Similarly, the sequence with the lowest associated weight at time $t = 5$ need not be the sequence with the lowest associated weight at time $t = 100$. The times at which all the quantiles converge correspond to rejuvenation times. For the good runs, the maximum weight is often much greater than that of the other quantiles. However, note that sometimes the 3rd quantile equals the maximum, though this characteristic is not apparent in these plots.

The distribution of the weights for a typical “bad” run displays a smaller variance than that of a typical “good” run since, in a good run, most sequences have very small weights. This phenomenon appears to conflict with the motivation for Rejuvenation, until we note that a set of “bad” imputed sequences may still have a high ESS. In fact, this plot suggests that in a bad run *all* the im-

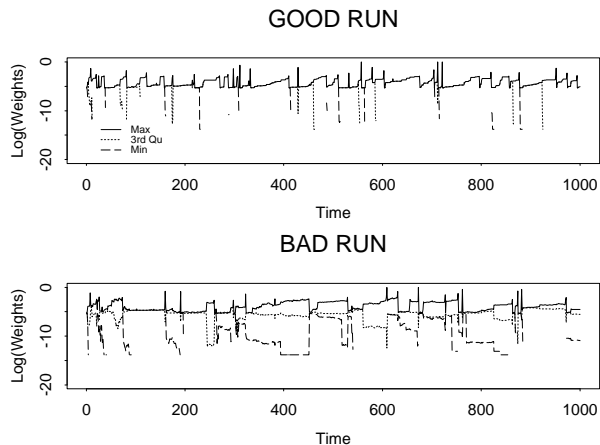


Figure 4: Looking at Distribution of Weights through a “good” run (BER=0.0) and “bad” run (BER=0.46). Channel: (0.407, 0.815, 0.407), $T=1000$, $N=200$, SNR=13db. There seems to be less variation in the distribution of the weights for “bad” runs.

puted sequences are poor so the sequence with maximum weight is simply the best sequence out of a set of bad ones. In this scenario, thresholding ESS will clearly not be sufficient for detecting when the procedure has gone amiss. Furthermore, rejuvenation will not help here because *all* the imputed sequences are bad.

2.3 Extensions to SIS

If there were some way to detect when the procedure had gone awry, then it would be possible to stop the imputation procedure and restart from time $t = 1$, with an entirely new set of imputed sequences. Liu et al. (1998) proposed some extensions to SIS which attempt to do just this. One such procedure is Rejection Control - SIS (RC-SIS). Figure 5 shows histograms of the error rates for $T = 100$, $N = 1000$ and $\sigma^2 = 0.05$ (i.e. SNR = 13db) when RC-SIS is applied to two lag-2 channels. The schedule of check points used for these simulations (and the corresponding thresholds as percentiles of the weights) was: {10 (80%), 20 (70), 30 (60), 40 (50), 50 (40), 60 (30), 70 (20), 80 (20), 90 (20), 99 (20)}.

Although this procedure seems to work well on the easy channel, the instability seen in Rejuvenation persists in RC-SIS when we consider a harder channel. There are a number of simulations with BER greater than 30% errors for SNR = 13db. The overall BER using RC-SIS was 0.0016 for the easy channel (0.946, 0.234, -0.167) and 0.0538 for the hard symmetric channel (0.407, 0.815, 0.407). These studies seem to suggest the following:

- Tracking ESS alone is not sufficient for detecting when the SIS has gone awry because ESS will be high if *all* imputed sequences are poor, but seems to work well if some of the sequences are good.




Figure 5: Histogram of Error Rate for Rejection Control applied to an Easy Channel (0.946, 0.234, -0.167), and a Hard Channel (0.407, 0.815, 0.407), $T=100$, $N=1000$, $SNR=13\text{db}$, based on 1000 simulation runs. Signs of instability persist in decoding the Hard Channel.

- Rejuvenation adds extra variation into the estimation procedure because sampling with replacement provides the bad sequences with an opportunity to have a greater impact on the final estimation of the signal.
- Perhaps the key lies in better initialization in either the channel or the imputed sequences.

In an effort to ameliorate the weight skewness problem, we investigated various ad-hoc extensions to the basic Rejuvenation procedure, partly motivated by Figures 3 and 4, including:

- Using the inter-quartile range instead of ESS, as a more robust measure of the distribution of the weights.
- Monitoring ESS and restarting SIS if there were signs of non-monotonic behavior between rejuvenation times.
- Partitioning imputed sequences into a small number of blocks. Running SIS within each block and then averaging across blocks to obtain the final sequence estimate.
- Running a Gibbs sampler, SIS, or LVCE, on the first few bits to obtain an initial estimate of the channel. Then, running Rejuvenation from time $t = 0$ using that channel estimate as the prior for the channel.
- Starting the procedure from all possible states (for small m) for the first few bits to guarantee that the correct sequence is in the initial set of imputed sequences.

Unfortunately, all of these schemes showed similar signs of instability. Even if we consider the slightly simpler

problem of using Rejuvenation in the presence of training data (e.g. first 10 bits being sent are known), the weight skewness problem persists. Least Squares was performed on the training data to obtain an initial estimate of the channel. Note that typically 2 to 3 times ($q + 1$) bits are required to obtain a reasonable estimate of the channel. However, the performance of these simulations was comparable to performance with no training data.

3 Discussion of Results

The comparisons made in Section 2 demonstrate that the Rejuvenation scheme is not very stable. It seems that “bad” sequences are occasionally sampled in the rejuvenating step, resulting in occasional simulations with an unusually large number of errors. Although these occasional bad simulations could be looked upon as outliers in the error rates, it is a statement in itself that such outliers crop up with this algorithm, but not in any of the other ones considered. The discrepancy between the heuristic argument for the Rejuvenation scheme’s performance, and its performance in practice, suggests that perhaps ESS is not the proper measure for deciding when to rejuvenate. Most of the extensions proposed by various authors depend on ESS, and consequently do not solve the weight skewness problem. Thus, it is questionable whether one would like to use SIS, or any of its variants, in a practical situation.

As a final comment we make note that the various SIS-based methods attempt to solve the fixed-parameter estimation problem which, perhaps, is not the correct forum for such methods. For the fixed-parameter problem, one cannot obtain uniform convergence results because the number of imputation streams has to increase exponentially over time to maintain a fixed precision. Consequently, it may not be advisable to use these algorithms for applications such as signal processing where the stream of data is very important.

References

- Ali, R., A. Murua, T. Richardson, and S. Roy (June 2001). A comparison of traditional methods and sequential bayesian methods for blind deconvolution problems. *Technical Report 397, Dept. of Statistics, University of Washington*. www.stat.washington.edu/www/research/reports/.
- Liu, J. and R. Chen (June 1995). Blind deconvolution via sequential imputations. *JASA 90, no. 430*, 567–576.
- Liu, J., R. Chen, and W. Wong (September 1998). Rejection control and sequential importance sampling. *JASA 93, no. 443*, 1022–1031.
- Seshadri, N. and C. Sundberg (February 1994). List viterbi decoding algorithms with applications. *IEEE Transactions and Communications*.