

MCMC METHODS FOR RESTORATION OF QUANTISED TIME SERIES

Paul T. Troughton and Simon J. Godsill

Signal Processing Group, Department of Engineering,
University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK
paul.troughton@ieee.org sjg@eng.cam.ac.uk

ABSTRACT

In digital systems, the amplitude of a time series is quantised with finite resolution. This is a nonlinear process which introduces distortion.

We develop a Bayesian, model-based approach to reducing the quantisation distortion when moving a time series, such as an audio signal, to a higher resolution medium. The signal is modelled as a discrete-time, continuous-valued autoregressive (AR) process of unknown order.

The model parameters and reconstructed signal are estimated using Markov chain Monte Carlo (MCMC) techniques. This requires samples to be drawn from a truncated multivariate Gaussian distribution, for which a Metropolis-Hastings approach is developed.

1. QUANTISATION DISTORTION

For digital processing, transmission, or storage, a signal is represented as discrete in both time (due to sampling) and value (due to quantisation). The quantisation process introduces an error component, often referred to as ‘quantisation noise’. Since this quantisation error is signal-dependent, it is perhaps better described as *quantisation distortion*.

The distortion consists of a series of added harmonics, together with aliased images of them. For complicated signals, this can be quite innocuous, behaving very much like white noise. When the signal is simple, however, the structure becomes clear. Furthermore, the relative levels of the added harmonics (and their inharmonic aliases) vary significantly with small changes in the level of the input signal. This behaviour can be disturbing in music signals, especially with those instruments, such as pianos, whose note waveforms become more sinusoidal as they decay. This effect produces what is known as ‘granulation noise’ [1].

Quantisation occurs both during analogue-to-digital conversion and during any subsequent manipulation of the digital signal that increases the word length, such as multiplication.




Figure 1: Modelling of quantised signal: e is the excitation process, x the undistorted signal, and y the observed, quantised signal.

Although there has been much research into the effects of quantisation and their elimination at the time of quantisation through the use of dither (see [2] for a comprehensive survey), there does not appear to be any published work on reconstructing an unknown signal, such as audio, when only a coarsely quantised version is available.

2. MODELLING FRAMEWORK

As shown in figure 1, we model the signal and the quantisation process explicitly. We assume stationarity of the process; nonstationary signals, such as audio, can be processed in short blocks.

2.1. Source model

We model the unquantised signal x as a (discrete-time) continuous-valued autoregressive process of order k with parameters $\mathbf{a}^{(k)}$, excited by white zero-mean Gaussian noise e with variance σ_e^2 . This can be represented in matrix-vector form as:

$$\mathbf{e} = \mathbf{A}^{(k)}\mathbf{x} = \mathbf{x}_1 - \mathbf{X}^{(k)}\mathbf{a}^{(k)} \quad p(\mathbf{e}) = \mathbf{N}(\mathbf{e} \mid \mathbf{0}, \sigma_e^2\mathbf{I})$$

where \mathbf{x}_0 and \mathbf{x}_1 are formed by partitioning \mathbf{x} into, respectively, the first k values and the remainder, and the matrices \mathbf{A} and $\mathbf{X}^{(k)}$ take appropriate forms [3]. The conditional likelihood for \mathbf{x} can then be expressed as:

$$p(\mathbf{x}_1 \mid k, \mathbf{a}^{(k)}, \sigma_e^2, \mathbf{x}_0) \approx \mathbf{N}(\mathbf{A}^{(k)}\mathbf{x} \mid \mathbf{0}, \sigma_e^2\mathbf{I}) \quad (1)$$

This work is supported by the Engineering & Physical Sciences Research Council, U.K.

2.2. Perfect quantiser

The quantisation process is deterministic:

$$\mathbf{y} = Q(\mathbf{x})$$

where \mathbf{y} is the observed, quantised signal and $Q(\cdot)$ is the quantisation function. For a perfect quantiser, $Q(x_t) = n\Delta$ if $(n - \frac{1}{2})\Delta \leq x_t < (n + \frac{1}{2})\Delta$, for integer n . If the quantiser's step height, Δ , is one then $Q(\cdot)$ is equivalent to rounding. In practical analogue-to-digital converters, quantisation is often far from perfect, possibly exhibiting uneven step heights and missed codes (see *e.g.* [4]). If these defects are memoryless, they can be incorporated straightforwardly into $Q(\cdot)$.

2.3. Prior distributions

Prior distributions represent our knowledge of the parameter values before analysing the signal. Since we are performing Bayesian model selection, the distributions must be proper. As in [5], k , $\mathbf{a}^{(k)}$, σ_a^2 and σ_e^2 are assumed to be *a priori* independent. We choose simple conjugate prior distributions which can be made quite uninformative with suitable hyperparameter values:

$$p(k) = \begin{cases} \frac{1}{k_{\max}+1} & k \in \{0, 1, \dots, k_{\max}\} \\ 0 & \text{elsewhere} \end{cases}$$

$$p(\mathbf{a}^{(k)} | k) = \mathbf{N}(\mathbf{a}^{(k)} | \mathbf{0}, \sigma_a^2 \mathbf{I}_k)$$

$$p(\sigma_a^2) = \text{IG}(\sigma_a^2 | \alpha_a, \beta_a)$$

$$p(\sigma_e^2) = \text{IG}(\sigma_e^2 | \alpha_e, \beta_e)$$

where $\mathbf{N}(\cdot)$ is the multivariate Gaussian distribution and $\text{IG}(\cdot)$ represents the inverse Gamma distribution, which can closely approximate the Jeffreys' prior.

3. BAYESIAN ESTIMATION & MCMC

We wish to reconstruct \mathbf{x} given only \mathbf{y} . Bayesian inference about \mathbf{x} is made on the basis of the marginal posterior distribution $p(\mathbf{x} | \mathbf{y})$. We cannot evaluate this analytically, so we use an MCMC approach (see *e.g.* [6]), in which we:

- Construct a Markov chain to draw (correlated) samples from the joint posterior, $p(\mathbf{x}, k, \mathbf{a}^{(k)}, \sigma_e^2, \sigma_a^2 | \mathbf{y})$
- Use these samples to calculate a Monte Carlo estimate of $p(\mathbf{x} | \mathbf{y})$.

We now look at the moves necessary to form the Markov chain.

3.1. AR order and parameters

As described in [5], we can sample $k, \mathbf{a}^{(k)} \sim p(k, \mathbf{a}^{(k)} | \mathbf{x}, \sigma_e^2, \sigma_a^2)$. Reversible-jump moves [7] are used because of

the variable dimensionality. Moves from a model of order k to one of order k' are proposed by sampling from a discretised Laplacian density centred on k :

$$k' \sim J(k \rightarrow k') \propto \exp(-\frac{1}{\lambda} |k' - k|)$$

Proposing the entire new parameter vector, $\mathbf{a}^{(k')}$, from its full conditional distribution, which is available analytically:

$$p(\mathbf{a}^{(k')} | k', \mathbf{x}, \sigma_a^2, \sigma_e^2) \propto \mathbf{N}(\mathbf{a}^{(k')} | \boldsymbol{\mu}_{\mathbf{sa}^{(k')}}, \mathbf{C}_{\mathbf{sa}^{(k')}})$$

where

$$\mathbf{C}_{\mathbf{sa}^{(k')}}^{-1} = \sigma_e^{-2} \mathbf{X}^{(k')T} \mathbf{X}^{(k')} + \sigma_a^{-2} \mathbf{I}_{k'}$$

$$\boldsymbol{\mu}_{\mathbf{sa}^{(k')}} = \sigma_e^{-2} \mathbf{C}_{\mathbf{sa}^{(k')}}^T \mathbf{X}^{(k')T} \mathbf{x}$$

leads to the acceptance probability for the model move being independent of both $\mathbf{a}^{(k)}$ and $\mathbf{a}^{(k')}$:

$$A((k, \mathbf{a}^{(k)}) \rightarrow (k', \mathbf{a}^{(k')}))$$

$$= \min \left(1, \frac{\sigma_a^{-k'} |\mathbf{C}_{\mathbf{sa}^{(k')}}|^{1/2} J(k' \rightarrow k)}{\sigma_a^{-k} |\mathbf{C}_{\mathbf{sa}^{(k)}}|^{1/2} J(k \rightarrow k')} \frac{\exp(\frac{1}{2} \boldsymbol{\mu}_{\mathbf{sa}^{(k')}}^T \mathbf{C}_{\mathbf{sa}^{(k')}}^{-1} \boldsymbol{\mu}_{\mathbf{sa}^{(k')}})}{\exp(\frac{1}{2} \boldsymbol{\mu}_{\mathbf{sa}^{(k)}}^T \mathbf{C}_{\mathbf{sa}^{(k)}}^{-1} \boldsymbol{\mu}_{\mathbf{sa}^{(k)}})} \right)$$

This approach gives a well mixed chain, with fast convergence [5]. The hyperparameters are sampled using straightforward Gibbs sampler moves.

3.2. Reconstructed signal

The quantisation process is a many-to-one mapping, so its inverse is one-to-many and provides the range of possible values that the input might have taken, *i.e.*

$$\mathbf{x} \in Q^{-1}(\mathbf{y})$$

Combining this knowledge with the conditional likelihood (eq. 1) for an autoregressive signal, gives:

$$p(\mathbf{x}_1 | \mathbf{y}, k, \mathbf{a}^{(k)}, \sigma_e^2, \mathbf{x}_0)$$

$$\propto \begin{cases} p(\mathbf{x}_1 | k, \mathbf{a}^{(k)}, \sigma_e^2, \mathbf{x}_0) & \mathbf{x} \in Q^{-1}(\mathbf{y}) \\ 0 & \text{elsewhere} \end{cases} \quad (2)$$

Since it is not computationally feasible to sample from the whole of \mathbf{x} directly (see §4), a Gibbs sampler approach is taken, in which a part of the block, \mathbf{x}_u , is sampled conditional on the remainder, \mathbf{x}_f . This is repeated, with different partitioning, until the whole block has been sampled.

The partitioning is as follows

$$\mathbf{e} = \mathbf{A}^{(k)} \mathbf{x} = \mathbf{A}_u^{(k)} \mathbf{x}_u + \mathbf{A}_f^{(k)} \mathbf{x}_f \quad (3)$$

where $(\cdot)_f$ also contains \mathbf{x}_0 and the first k samples from the following block, to ensure continuity across the block boundaries.¹ $(\cdot)_u$ need not be contiguous, but since it is likely that adjacent samples will be more highly correlated, it may be advantageous to sample them jointly.

The full conditional distribution for \mathbf{x}_u is the same as that required for interpolation [8, §12.5.1], except that it is bounded to lie within $Q^{-1}(\mathbf{y}_u)$. If the bounds are represented by the function

$$B(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in Q^{-1}(\mathbf{y}) \\ 0 & \text{elsewhere} \end{cases}$$

then equation (2) can be rearranged as

$$\begin{aligned} p(\mathbf{x}_u | \mathbf{y}, \mathbf{x}_f, k, \mathbf{a}^{(k)}, \sigma_e^2) \\ \propto B(\mathbf{x}_u) \times p_e(\mathbf{A}_u^{(k)} \mathbf{x}_u + \mathbf{A}_f^{(k)} \mathbf{x}_f) \\ \propto B(\mathbf{x}_u) \times \mathbf{N}(\mathbf{A}_u^{(k)} \mathbf{x}_u | -\mathbf{A}_f^{(k)} \mathbf{x}_f, \sigma_e^2) \\ \propto B(\mathbf{x}_u) \times \mathbf{N}(\mathbf{x}_u | \boldsymbol{\mu}_{l\mathbf{x}_u}, \mathbf{C}_{l\mathbf{x}_u}) \end{aligned} \quad (4)$$

where

$$\begin{aligned} \mathbf{C}_{l\mathbf{x}_u} &= \sigma_e^2 (\mathbf{A}_u^{(k)T} \mathbf{A}_u^{(k)})^{-1} \\ \boldsymbol{\mu}_{l\mathbf{x}_u} &= -(\mathbf{A}_u^{(k)T} \mathbf{A}_u^{(k)})^{-1} \mathbf{A}_u^{(k)T} \mathbf{A}_f^{(k)} \mathbf{x}_f \end{aligned} \quad (5)$$

which is a multivariate Gaussian distribution, truncated to a hypercube² in \mathbf{x}_u -space.

4. SAMPLING TRUNCATED GAUSSIANS

The usual method for sampling from an n -dimensional multivariate Gaussian distribution is to find, using the Cholesky decomposition, the linear transformation required to diagonalise the covariance matrix. A vector of n independent samples can then be drawn from a univariate Gaussian distribution and inverse transformed to give a sample from the multivariate distribution.

To obtain samples from a distribution $B(\boldsymbol{\theta}) \times \mathbf{N}(\boldsymbol{\theta} | \boldsymbol{\mu}, \mathbf{C})$, the simplest method would simply be to draw samples from the unbounded Gaussian and reject those which lie outside the bounded area. The average acceptance rate will equal the proportion of the probability mass of the Gaussian distribution which lies within the bounds. It will hence tend to decrease as the number of dimensions increases. If the bounds are both out in one tail of the distribution, it could be vanishingly small.

Efficient methods have been developed for sampling from univariate truncated Gaussian distributions by rejection sampling using different Rayleigh or exponential distributions depending on the positions of the bounds [9, 10] or by using the inverse cumulative distribution function [11].

¹In fact the structure of $\boldsymbol{\mu}_{l\mathbf{x}_u}$ (eq. 5) is such that the $(\cdot)_f$ partition need only contain the nearest k samples to each side of the $(\cdot)_u$ partition.

²This is for a perfect quantiser; more generally it will be a cuboid region.

In this case, it is not possible simply to draw from univariate truncated distributions and transform: although the faces of the cuboid forming the bounds are parallel to the axes in $\boldsymbol{\theta}$ -space, they generally will not be after transformation, so the bounds on each element will depend on the values of the others, preventing independent sampling.


This problem can be addressed, at the cost of producing dependent samples, using MCMC methods.

4.1. Gibbs sampling

One way to produce samples from the truncated distribution is to use a Gibbs sampler, in which each element is sampled from its full conditional distribution, in this case a truncated univariate Gaussian. This method is suggested in [10] and [11]. For this problem, it can be implemented straightforwardly by making the partition $(\cdot)_u$ in equation (3) contain a single element. Equation (4) then becomes the required full conditional distribution.

This is a potential drawback with this method: in an MCMC scheme, when some variables are highly correlated, convergence will tend to be slow unless those variables are sampled jointly [12]. Hence, if the non-diagonal elements of \mathbf{C} are significant, convergence may be slow.

The inset probability contour plot illustrates the problem in two dimensions: the sampler can only move parallel to the axes, and will tend to stay in relatively high probability regions, so it cannot



move between points A and B quickly. In higher dimensions, this lack of mobility becomes serious.

4.2. Gaussian windowing

In order to be able to sample multiple components jointly, we now describe an alternative technique, which does not appear to have been suggested elsewhere.

As discussed earlier, direct rejection sampling using the unbounded distribution is very inefficient if most of the probability lies outside the bounds. Multiplying the (unbounded) target Gaussian distribution by another multivariate Gaussian, centred within the bounds, results in a distribution which is related to the target distribution but has a much greater probability within the bounds. The combined

distribution is itself a multivariate Gaussian:

$$\underbrace{\mathbf{N}(\boldsymbol{\theta} \mid \boldsymbol{\mu}_c, \mathbf{C}_c)}_{\text{Combined}} \propto \underbrace{\mathbf{N}(\boldsymbol{\theta} \mid \boldsymbol{\mu}, \mathbf{C})}_{\text{Target}} \times \underbrace{\mathbf{N}\left(\boldsymbol{\theta} \mid \frac{\mathbf{b}_{\max} + \mathbf{b}_{\min}}{2}, \kappa \times \text{diag}(\mathbf{b}_{\max} - \mathbf{b}_{\min})\right)}_{\text{Window}}$$

where \mathbf{b}_{\max} and \mathbf{b}_{\min} are vectors containing the bounds on each element, κ is the *window factor* and \mathbf{C}_c and $\boldsymbol{\mu}_c$ can be derived straightforwardly.

Independent samples can be drawn from this combined distribution and rejection sampling can be used to enforce the bounds. The bias introduced by the windowing function can then be removed by an independence sampler step [6]:

$$\alpha(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') = \min \left(1, \frac{\mathbf{N}(\boldsymbol{\theta}' \mid \boldsymbol{\mu}, \mathbf{C})}{\mathbf{N}(\boldsymbol{\theta} \mid \boldsymbol{\mu}, \mathbf{C})} \frac{\mathbf{N}(\boldsymbol{\theta} \mid \boldsymbol{\mu}_c, \mathbf{C}_c)}{\mathbf{N}(\boldsymbol{\theta}' \mid \boldsymbol{\mu}_c, \mathbf{C}_c)} \right)$$

which can be simplified through cancellation.

The value of κ controls the trade-off between the acceptance rates in the rejection sampling step and in the independence sampler.

The acceptance rate falls as the dimension of $\boldsymbol{\theta}$ increases, so for a given acceptance rate, the number of elements of \mathbf{x} which can be sampled jointly will be limited. For the quantisation removal problem, experimentation has shown that joint sampling of 5 elements is acceptably efficient.

The structure of the covariance matrix (eq. 4) is such that adjacent samples will have highest correlation, so sampling subblocks \mathbf{x}_u of consecutive samples should give fastest convergence. If fixed subblocks were used, samples on the subblock boundaries would never be sampled jointly with those on the adjacent subblock. This problem is avoided by applying a random offset to the subblocking on each iteration.

4.3. Comparison for synthetic AR data

The performance of the two methods was compared for a simple quantisation distortion problem. An AR(2) signal was synthesised, with poles at $0.99e^{j\pi \frac{\pm 15}{180}}$ and white Gaussian excitation of unit variance. The signal was quantised with a step height of 20, which introduced quantisation distortion 9.5 dB (rms) below the signal. A single block of 512 samples was restored, with the AR model order fixed at two.³ Fifty runs of 100 iterations were made with each of the algorithms. In the windowing algorithm, κ was 10 and blocks of five consecutive samples were sampled jointly.

Figure 2 shows the distortion:signal ratio at each iteration, averaged across the 50 runs of each algorithm. It can

³To avoid end effects, which become insignificant in signals with many blocks, two samples before and after the block were made available from the unquantised signal.




Figure 2: Comparison of truncated Gaussian sampling algorithms: (*dotted*) univariate Gibbs sampler, and (*solid*) windowed Gaussian method. The plot shows performance in reducing quantisation distortion in a synthetic AR(2) signal.

be seen that the windowing algorithm exhibits much faster convergence. The total computation time was also significantly less than for the Gibbs sampling algorithm.

4.4. Discussion

On the basis of this test, the joint sampling algorithm is chosen. The efficiency might be improved further by changing the windowing function. A single Gaussian is not very close in shape to the required rectangular window. A closer fit could be achieved by using a mixture of Gaussians, for example equispaced Gaussians across the bounded region. A smaller value of κ could then be used, lowering the rejection rate. The drawback is the additional complexity involved in sampling from the mixture and evaluating its density function for use in calculating the acceptance probability.

5. RESULTS

A short extract of low level, decaying notes was taken from a commercial 44.1 kHz, 16 bit recording of a piano piece. In the original recording ($\Delta = 1$), no distortion was audible; requantising the signal with $\Delta = 50$ introduced distortion 23 dB (rms) below the signal, which could be clearly heard.

The signal was split into blocks of 1024 samples, and the sampler run for 200 iterations. Initial values were $\mathbf{x} = \mathbf{y}$ and $k = 6$; the remainder of the parameters were drawn from their full conditionals. k_{\max} was set to 50. A Monte Carlo estimate, $\hat{\mathbf{x}}$, of the signal was made using the final 100 iterations, in which the distortion was reduced by an average of 10 dB (rms), a clearly audible improvement.

Figure 3 shows the signal together with, for each block, estimates of $p(k)$ and $p(\sigma_e)$ and the distortion level before




Figure 3: Quantised piano signal: signal; histograms of estimated marginal posterior distributions (in which darkness represents probability); and distortion levels.




Figure 4: Part of block 80: (top) signal and (bottom) error, in both of which the faint line corresponds to the quantised signal and the heavy line to the restored one.

and after restoration. Figure 4 shows part of the signal in which the quantisation distortion was audible. It can be seen that error signal is much smaller after restoration. For the part shown, which is typical of the quieter blocks, a 15 dB (rms) improvement was achieved.

Model orders $2 \leq k \leq 16$ were observed in the sampler output. Similar experiments using fixed model orders showed the importance of model selection: if the model order was 2, the improvement was only 7.5 dB (rms) and noise could be heard in the restored signal; if it was fixed at 30, there were disturbing ‘musical noise’ artifacts [8, §6.1.2] where some distortion elements were being modelled as signal.

6. DISCUSSION

We have developed a technique for fitting an AR model to a signal of which only a quantised version is available. This required the development of a method for sampling from

truncated multivariate Gaussian distributions.

Audio signals tend to exhibit a $1/f$ spectral shape. In order to prevent spurious poles at high frequencies, which have been observed, future work will investigate priors on $\mathbf{a}^{(k)}$ which incorporate this knowledge,

The signals which tend to cause granulation noise are those which have strong sinusoidal components. In initial experiments, sinusoid + white noise models have not performed well, but adding sinusoidal basis functions to an AR model has improved performance.

7. REFERENCES

- [1] R. C. Maher. “On the nature of granulation noise in uniform quantization systems”. *Journal of the Audio Engineering Society*, 40(1–2):12–20, 1992.
- [2] S. P. Lipshitz, R. A. Wannamaker, & J. Vanderkooy. “Quantization and dither: a theoretical survey”. *Journal of the Audio Engineering Society*, 40(5):355–375, 1992.
- [3] G. E. P. Box, G. M. Jenkins, & G. C. Reinsel. *Time Series Analysis: Forecasting and Control*. Holden-Day, 3rd edition, 1994.
- [4] M. Tsukamoto, K. Matsunaga, O. Morioka, T. Saito, T. Igarashi, H. Yazawa, & Y. Takahashi. “Correction of nonlinearity errors contained in the digital audio signals”. Presented at the 104th Convention of the Audio Engineering Society, preprint 4698, 1998.
- [5] P. T. Troughton & S. J. Godsill. “A reversible jump sampler for autoregressive time series”. *Proceedings of IEEE ICASSP-98*, IV:2257–2260, 1998.
- [6] W. R. Gilks, S. Richardson, & D. J. Spiegelhalter, eds. *Markov Chain Monte Carlo in Practice: Interdisciplinary Statistics*. Chapman & Hall, 1996.
- [7] P. J. Green. “Reversible jump Markov chain Monte Carlo computation and Bayesian model determination”. *Biometrika*, 82(4):711–732, 1995.
- [8] S. J. Godsill & P. J. W. Rayner. *Digital Audio Restoration: A Statistical Model Based Approach*. Springer-Verlag, 1998.
- [9] G. Marsaglia. “Generating a variable from the tail of the normal distribution”. *Technometrics*, 6(1):101–102, 1964.
- [10] C. P. Robert. “Simulation of truncated normal variables”. *Statistics and Computing*, 5:121–125, 1995.
- [11] J. H. Kotecha & P. M. Djurić. “Gibbs sampling approach for generation of multivariate Gaussian random variables”. *Proceedings of IEEE ICASSP-99*, 3:1757–1760, 1999.
- [12] J. S. Liu, W. H. Wong, & A. Kong. “Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes”. *Biometrika*, 81(1):27–40, 1994.