Optimal quantization and suprathreshold stochastic resonance

Mark D. McDonnell*[a], Nigel G. Stocks[b], Charles E.M. Pearce[c] and Derek Abbott[a],

[a]Centre for Biomedical Engineering (CBME) and
School of Electrical & Electronic Engineering,
The University of Adelaide, SA 5005, Australia

[b]School of Engineering,
The University of Warwick, Coventry CV4 7AL, UK

[c]School of Mathematical Sciences
The University of Adelaide, SA 5005, Australia

ABSTRACT

It is shown that Suprathreshold Stochastic Resonance (SSR) is effectively a way of using noise to perform quantization or lossy signal compression with a population of identical threshold-based devices. Quantization of an analog signal is a fundamental requirement for its efficient storage or compression in a digital system. This process will always result in a loss of quality, known as distortion, in a reproduction of the original signal. The distortion can be decreased by increasing the number of states available for encoding the signal (measured by the rate, or mutual information). Hence, designing a quantizer requires a tradeoff between distortion and rate. Quantization theory has recently been applied to the analysis of neural coding and here we examine the possibility that SSR is a possible mechanism used by populations of sensory neurons to quantize signals. In particular, we analyze the rate-distortion performance of SSR for a range of input SNR's and show that both the optimal distortion and optimal rate occurs for an input SNR of about 0 dB, which is a biologically plausible situation. Furthermore, we relax the constraint that all thresholds are identical, and find the optimal threshold values for a range of input SNRs. We find that for sufficiently small input SNRs, the optimal quantizer is one in which all thresholds are identical, that is, the SSR situation is optimal in this case.

Keywords: suprathreshold stochastic resonance, stochastic resonance, quantization, rate-distortion, Analog to Digital Conversion, noise, neural coding, dithering, lossy source coding, compression

1. INTRODUCTION

Stochastic resonance is the term used to described systems in which the presence of input or internal noise provides the optimal output of that system.[1–3] Although the term was originally used to refer to the very specific case of nonlinear systems driven by periodic input signals, subject to additive white noise, and performance measured by the output signal to noise ratio, the name is now applied very broadly to any system in which some nonzero level of noise can provide a performance improvement. Examples of systems in which SR has been described include a Schmitt Trigger circuit,[4] ring lasers,[5] neurons,[6] SQUIDS[7] and ion channels.[8] The literature on SR also contains many studies of systems consisting of a single threshold device (see for example[9–11]). SR occurs in such systems when the addition of noise to a subthreshold signal provides a system output when no output would have occurred otherwise.

By contrast, the term Suprathreshold Stochastic Resonance (SSR) was coined[12] for the multi-threshold system is shown in Figure 1. For SSR, the system consists of $N > 1$ thresholds, with identical values, but subject to *iid* additive threshold noise in which case SR effects occur regardless of whether the input signal is sub or supra threshold.[12, 13] In contrast to dithering in the process of quantization,[14] all thresholds were set to the same value, and the noise signal was allowed to be very large and independent on each threshold. In keeping with SR conventions, the system was analyzed by calculating the variation in a measure of system performance as the internal noise level increases from zero. Due to the system input being *iid* samples from a stationary random

*mmcdonne@eleceng.adelaide.edu.au; phone +61 8 8303 6296; fax +61 8 8303 4360

Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems III,
edited by Nigel G. Stocks, Derek Abbott, Robert P. Morse, Proc. of SPIE Vol. 5841
(SPIE, Bellingham, WA, 2005) · 0277-786X/05/$15 · doi: 10.1117/12.609542

signal, mutual information was a natural measure to use. It was shown numerically that the maximum mutual information through the system occurred for the thresholds set equal to the signal mean, and some nonzero value of the ratio of noise variance, σ_η^2, to signal variance, σ_x^2. Label this ratio as σ and note that it is effectively the reciprocal of the input Signal to Noise Ratio (SNR).

Part of the initial motivation for studying the SSR model was its similarity to populations of neurons. The essential nonlinearity in most well-known neuron models is that of a threshold which emits output spikes when crossed by an input signal. Since its initial study, the SSR effect has also been shown to occur in populations of FitzHugh-Nagumo neuron models[15] and applied to cochlear implant encoding.[16–18]

The reason that SSR occurs is that in the absence of noise, only two output states are possible— the case of all thresholds "on" or all "off," in which case the mutual information is at most 1 bit per sample. For a small amount of noise all $N + 1$ output states can occur, but most only with a very small probability. For the optimal amount of noise, the output states are each occupied with probabilities that on average best reflect the input signal distribution. For very large noise, the relationship between the input and output distributions become very close to independent, and very little information is transmitted.

Given that this implies that the SSR situation is sub-optimal in the absence of noise (since $\log_2(N + 1)$ bits per sample are available[12]), the obvious question is to ask whether there are any values of input SNR for which SSR is optimal. The answer to this is yes. We have previously shown that SSR becomes the optimal threshold configuration for input SNR's of the order of 0 dB.[19] This is a biologically plausible situation for sensory neurons.[20]

Since the SSR model converts a continuously valued input signal into a discretely valued output signal, it is effectively a stochastic quantization scheme. Hence, it can also be analyzed in terms of standard quantization theory. This forms the first part of the goals of this paper. Secondly, we will relax the SSR constraint that all thresholds are equal to the same value, and show that SSR is the optimal quantization for input SNR's of about 0 dB or smaller. This will be followed by analyzing this situation in terms of quantization theory, and a comparison with SSR. We begin however, by formulating a mathematical description of the general model of Figure 1.

2. GENERAL MODEL

The model we use is shown in Figure 1. Assume a stationary input signal with PDF $P_x(x)$ and variance σ_x^2. The input to the multi-threshold system is assumed to consist of independent samples from this signal. All N threshold devices receive the same signal sample, x. Denote the n-th threshold value as θ_n. All thresholds are subject to *iid* additive input noise, η_n, with PDF $R(\eta)$. Throughout this work we assume that both $P_x(x)$ and $R(\eta)$ are even valued functions about a mean of zero, and have infinite support. Label the output of each threshold element as y_n, where $y_n = 0.5\text{sign}(x + \eta_n - \theta_n) + 0.5$ so that $y_n \in \{0, 1\}$. Hence, y_n is equal to unity, or "on," if the signal plus noise is greater than the n-th threshold, θ_n, and zero, or "off," otherwise. The overall output of the quantizer is then $y = \sum_{n=0}^{N} y_n$. Thus, y is a discretely valued stochastic encoding of x, such that $y \in \{0, .., N\}$.

Measures of performance will depend on calculating the joint input-output PDF, $P_{xy}(x, n) = P(n|x)P_x(x)$, where $P(n|x)$ is the probability that the output y is equal to n given input sample x. Integration with respect to x gives the probability that y is equal to n, $P_y(n) = \int_{-\infty}^{\infty} P(n|x)P_x(x)dx$. Hence, if $P_x(x)$ and $P(n|x)$ are known, then all other distributions can be derived, as can the performance measures we are interested in, as shown in the subsequent sub-sections.

We will also make extensive use of the probability that the n-th device is "on," given x, that is, the probability that $x + \eta_n \geq \theta_n$. Let this be $\hat{P}_n(x)$. If the noise has PDF $R(\eta)$ then

$$\hat{P}_n(x) = 1 - \int_{\eta=-\infty}^{\eta=\theta_n - x} R(\eta)d\eta = 1 - F_R(\theta_n - x), \tag{1}$$

where F_R is the cumulative distribution function of the noise. If $R(\eta)$ is even about a mean of zero then

$$\hat{P}_n(x) = F_R(x - \theta_n). \tag{2}$$

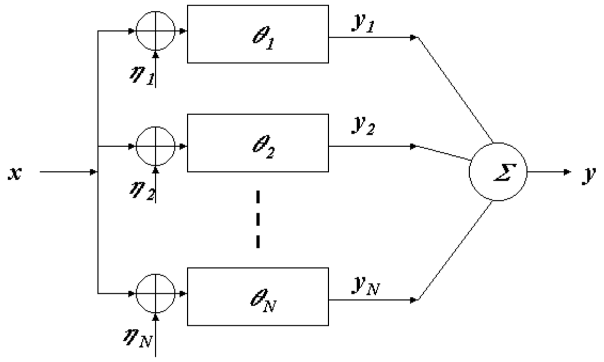

Figure 1. Model of the stochastic quantizer. A sample, x, from the source distribution, $P_x(x)$, is quantized by N threshold elements, all of which are subject to *iid* additive noise, η_n. The output from the n-th threshold, y_n, is unity if the sum of x and η_n is greater than the threshold value, θ_n, and zero otherwise. The overall output, y, is the sum of the y_n's.

We will now consider two measurers of performance: the mutual information, and the mean square distortion.

2.1. Mutual information

The mutual information for the multi-threshold system is that of a semi-continuous channel,[12]

$$I(x,y) = -\sum_{n=0}^{N} P_y(n)\log_2 P_y(n) - \left(-\int_{-\infty}^{\infty} P_x(x) \sum_{n=0}^{N} P(n|x)\log_2 P(n|x)dx \right).$$

This can be calculated by numerical integration for any N assuming $P_x(x)$, and $P(n|x)$ are known.

2.2. Mean Square Distortion

The array of comparators can be considered as a stochastic quantizer. Quantization theory[21] generally models a quantizer as having two components. The first is an encoding component, which assigns some output codeword to a given range of values of the input signal. Most quantization theory considers only deterministic encoding. Here, the encoding is probabilistic, so that a given input value, x, may be assigned to any of the output codewords, $y \in [0, N]$ with probability $P(n|x)$. The second part of the generic quantizer model is the decoding. The decoder assigns a value to each codeword at the output of the encoder that allows an approximate reconstruction of the input signal to be made. Generally, quantizer performance is considered in terms of some distortion measure, the most common of which is the mean square error between the original signal and the reconstructed signal.

Let the output of the array of comparators be decoded so that each value of $y \in [0,..,N]$ is decoded to $\hat{y} \in [\hat{y}_0..,\hat{y}_N]$, where $\hat{y}_n$ is the "reproduction point" for encoder value n. Thus, the Mean Square Error (MSE) distortion is

$$D = \int_x \sum_{n=0}^{N} (x-\hat{y})^2 P(n|x) P_x(x) dx, \tag{3}$$

which can also be calculated for any N and $\hat{y}$, if $P_x(x)$ and $P(n|x)$ are known.

An output SNR measure— often known in quantization theory as the Signal to Quantization Noise Ratio (SQNR)— can be defined, in terms of the ratio of the input signal variance to the MSE distortion as

$$\text{SQNR} = 10\log_{10}\left(\frac{\sigma_x^2}{D}\right) \quad \text{dB}. \tag{4}$$

It is possible to show that an optimal mean-square error decoding exists. Labelling this decoding as $\hat{x}$, such a decoding is the expected value of x given y,[22]

$$\hat{x}_n = \mathrm{E}[x|n] = \frac{\int_x xP(x)P(n|x)dx}{\int_x P(x)P(n|x)dx}. \tag{5}$$

Using this decoding, the Minimum Mean Square Error (MMSE) distortion is

$$\hat{D} = \mathrm{E}[x^2] - \mathrm{E}[\hat{x}_n^2]. \tag{6}$$

Noting that $\mathrm{E}[x] = \mathrm{E}[\hat{x}] = 0$, the MMSE is the difference between the variance of the input signal and the variance of the optimally decoded output signal.

2.3. Rate Distortion trade-off

In general, design of a good quantizer for a given source involves a trade-off between rate and distortion. That is, decreasing the average distortion will always require an increase in the rate required. This trade-off is measured using the rate-distortion function,[23] often expressed as $R(D)$, where R is the rate, and D is the distortion measure. Shannon proved that a lower bound exists for $R(D)$ for a Gaussian source with power σ_x^2, and the mean square distortion measure,[23] and that this bound is given by

$$R(D) = 0.5\log_2\left(\frac{\sigma_x^2}{D}\right). \tag{7}$$

This equation says that no quantization scheme can achieve a distortion less than D with a rate lower than $R(D)$ for a Gaussian source. Or in other words, a quantization scheme with rate R will provide a mean square distortion no smaller than D. Note that (7) can be rearranged to give a lower bound on the SQNR in decibels of about $6.02R$. This is a well-known rule of thumb in quantizer design that states that a one bit increase in the number of output bits gives about a 6 dB increase in SNR.

3. SUPRATHRESHOLD STOCHASTIC RESONANCE

Consider the specific case of all thresholds set equal to the same value, θ. As mentioned in the introduction, this is the specific case of the SSR model. It has been shown that the mutual information between the input and output is optimized for SSR when θ is equal to the signal mean.[13] We will only consider signal and noise distributions with PDFs that are even about a mean of zero. Hence, let $\theta = 0$ for all n to give the maximum information SSR situation. Let $\hat{P}(x)$ be the probability that a device with a threshold of zero is "on", for a given x. Thus from (2), $\hat{P}(x) = F_R(x)$.

The conditional probability mass function of y given x is given by the binomial formula,[12]

$$P(n|x) = \binom{N}{n}\hat{P}(x)^n(1-\hat{P}(x))^{(N-n)}. \tag{8}$$

Using this formula, the mutual information simplifies to[12]

$$\begin{aligned} I(x,y) = &-\sum_{n=0}^{N} P_y(n)\log_2 P^*(n) \\ &+ N\int_x P(x)\left(\hat{P}(x)\log_2\hat{P}(x) + (1-\hat{P}(x)\log_2(1-\hat{P}(x))\right)dx, \end{aligned} \tag{9}$$

where $P^*(n) = P_y(n)/\binom{N}{n}$. Thus $P^*(n) = \int_x P(x)P_{1|x}^n(1-P_{1|x})^{N-n}dx$. The mutual information can therefore be calculated from knowledge of $P_x(x)$, $\hat{P}(x)$ and N. A plot of the mutual information against σ is shown in Figure 2 for various N and Gaussian signal and noise.

The mean square error decoding of the output of the SSR model was first examined by McDonnell *et al.*[24] A linearly spaced decoding was used to show that SSR effects also occur in the mean square distortion measure. For the case of Gaussian signal and noise distributions, the optimal linear decoding was calculated. Such a decoding was shown to vary with σ. Although not mentioned in this reference,[24] such an optimal linear decoding was later recognized as *Wiener decoding*.[25] Here, however, we use the optimal MSE distortion, as given by (6). A plot of MMSE against σ for various N and Gaussian signal and noise is shown in Figure 3.

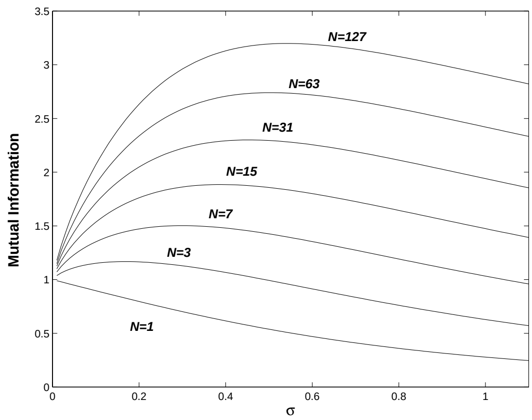

Figure 2. Plot of mutual information against σ for various N and a zero mean Gaussian source with unity variance, and Gaussian noise.

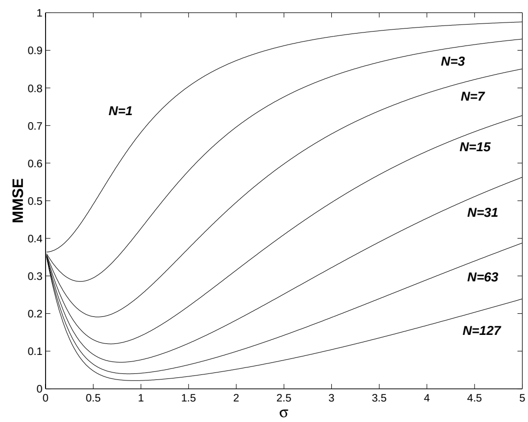

Figure 3. Plot of MMSE against σ for various N, and a zero mean Gaussian source with unity variance, and Gaussian noise.

3.1. Rate-distortion trade-off for SSR

The operational rate-distortion plot describes the actual trade-off between rate and distortion. This can be obtained for SSR simply by plotting the mutual information against the MMSE distortion as σ varies. This is shown in Figure 4 for zero mean Gaussian signal and noise and various N. The thick solid line in Figure 4 shows the theoretical $R(D)$ curve. Observe that the plot for a single value of N starts at a rate of one bit per sample at $\sigma = 0$, then increases with rate and decreases with MMSE, as σ increases. The rate then reaches its maximum before the MMSE reaches its minimum. Then with continuing increasing σ, the curve reaches its MMSE minimum, before curling back down towards the $R(D)$ curve. Note that this means that (except for very large σ) there are two values of σ for which the same distortion can occur, corresponding to two different rates. If the main goal of a quantizer is to operate with maximum output SNR, this observation indicates that the optimal value of σ to use is the one which achieves the minimum distortion, rather than the maximum rate. A further observation is the fact that for very large σ, the SSR quantization scheme provides an output that is very close to the theoretical rate-distortion curve.

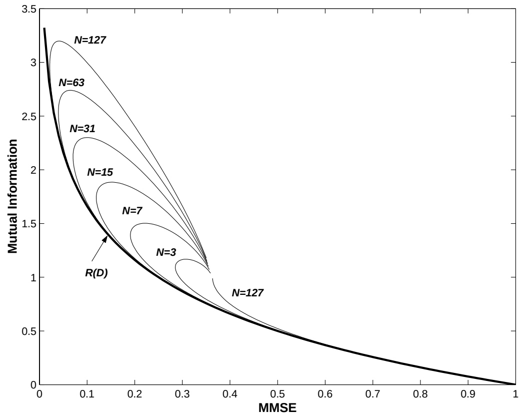

Figure 4. Plot of mutual information against MMSE for a zero mean Gaussian source with unity variance, Gaussian noise, and a number of values of N. The thick solid line shows $R(D)$. Note that there are in general two values of mutual information that achieve the same MMSE and two values of MMSE that achieve the same mutual information. This is due to the stochastic resonance behavior of the mutual information and MMSE for increasing σ.

4. ARBITRARY THRESHOLDS: OPTIMAL QUANTIZATION

4.1. Finding the optimal thresholds

We now aim to find the threshold settings that maximize the mutual information, or minimize the mean square distortion, as the input SNR varies. Such goals can be formulated as the following two nonlinear optimization problems:

$$\begin{aligned} \text{Find:} \quad & \max_{\{\theta_n\}} I(x, y) \\ \text{subject to:} \quad & \{\theta_n\} \in \mathbf{R^N}, \end{aligned} \tag{10}$$

or

$$\begin{aligned} \text{Find:} \quad & \min_{\{\theta_n\}} \text{MMSE} \\ \text{subject to:} \quad & \{\theta_n\} \in \mathbf{R^N}. \end{aligned} \tag{11}$$

It is necessary to solve (10) or (11) numerically using standard unconstrained non-linear optimization methods such as the conjugate gradient method.[26] Such methods require knowledge of $P(n|x)$ to calculate the mutual information or distortion. For arbitrary thresholds, it is very difficult to obtain such exact expressions however, given any arbitrary N, $R(\eta)$, and $\{\theta_n\}$, $\{\hat{P}_n(x)\}$ can be calculated exactly for any value of x from (1), from which $P(n|x)$ can be found numerically using an efficient recursive formula.[24]

Note that the objective function is not convex, and there exist a number of local optima. This problem can be overcome by employing random search techniques such as simulated annealing. We present here results for the optimal quantization obtained by solving (10) and (11) for Gaussian signal and Gaussian noise, $\sigma_x = 1$ and $N = 7$.

Firstly, Figures 5 and 6 show the mutual information and MMSE obtained by either maximizing mutual information or minimizing MMSE. Clearly, maximizing mutual information does not give the same result as minimizing MMSE, which illustrates the difference between these measures; mutual information takes into account other moments of the input-output distribution, not just the second moment. These figures also show the mutual information and MMSE obtained without optimizing the thresholds, that is, the SSR case. For maximized mutual information, the mutual information strictly decreases with increasing σ. For minimized MMSE, the MMSE strictly increases with increasing σ.

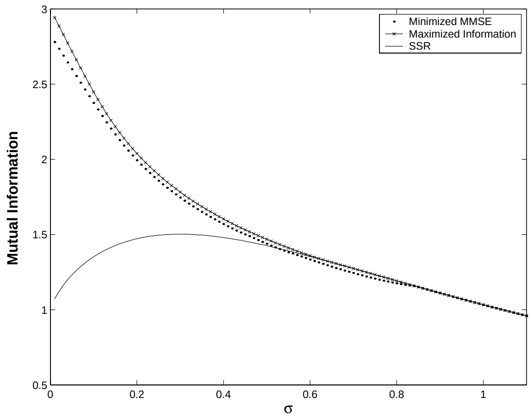

Figure 5. Plot of mutual information, against increasing σ, for $N = 7$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

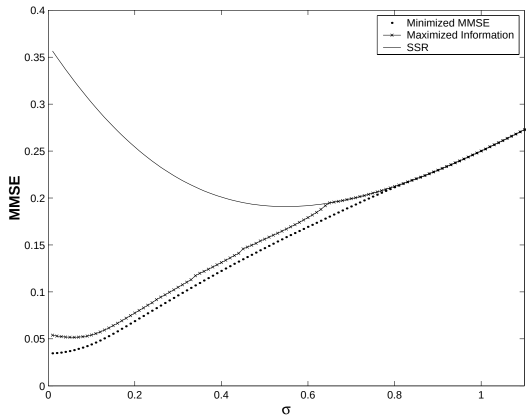

Figure 6. Plot of MMSE, against increasing σ, for $N = 7$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

Figure 7 shows the optimal thresholds for (10), and Figure 8 shows the optimal thresholds for (11).

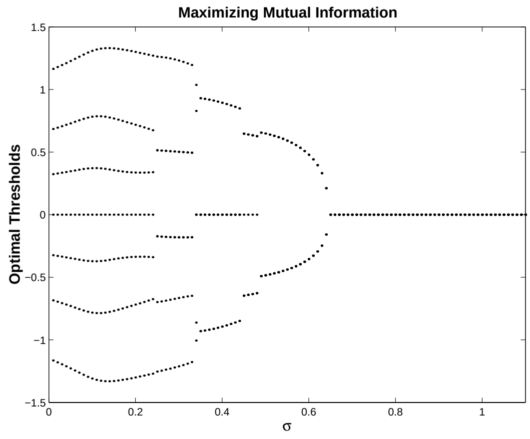

Figure 7. Plot of thresholds that maximize the mutual information, against increasing σ, for $N = 7$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

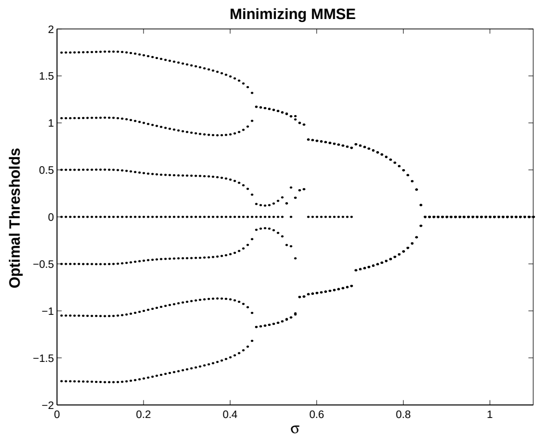

Figure 8. Plot of thresholds that minimize the MMSE, against increasing σ, for $N = 7$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

There are several notable features in these two plots, which both show the same qualitative behavior. Firstly, for small σ, the optimal thresholds are all uniquely valued, and widely distributed across the input distribution's support. For large σ, the optimal thresholds are all equal to the same value of zero. This is precisely the situation that we impose in the SSR model. This shows that the SSR situation of all thresholds identical is an optimal quantization for sufficiently low input SNR's, when there is *iid* additive threshold noise. For the mutual information, SSR is optimal at about $\sigma = 0.65$ and for MMSE, SSR becomes optimal at $\sigma = 0.85$. A further observation is the existence of bifurcations in the optimal threshold settings. Note that as σ increases, there are values of σ at which the number of unique threshold values decreases. This occurs several times with increasing σ until the point where SSR becomes optimal. We have found such behavior to persist for larger N, for other source and noise distributions, and other measures.[19]

4.2. Rate-distortion trade-off for arbitrary thresholds

The operational rate-distortion trade-off for the case of optimized thresholds is shown in Figure 9. This plot clearly shows that when the MMSE is minimized, the corresponding mutual information is smaller than that obtained by maximizing the mutual information. Conversely, maximizing the mutual information results in a larger distortion. When the distortion is minimized, the resultant curve is relatively close to the $R(D)$ curve, but as expected, does not reach it. This plot also clearly illustrates that the performance of SSR at best reaches about half the optimal mutual information, but only about five times the minimum possible distortion.

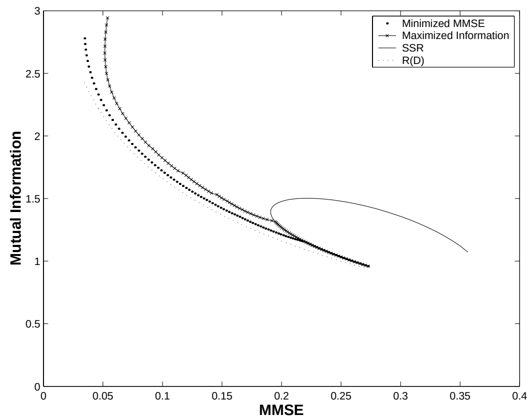

Figure 9. Plot of mutual information against MMSE for increasing σ, $N = 7$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

5. CONCLUSIONS

We have shown how SSR can be interpreted as a stochastic quantization scheme and analyzed it in terms of quantization theory. Furthermore, we have have examined optimization of the threshold values in this quantization scheme for a range of noise intensities. The most important result is that for sufficiently small input SNR's, the optimal quantization is the SSR situation of all thresholds equal to the signal mean. It is very important to note though that this does not apply if all thresholds have the same input noise. It is the independence of the noise at each threshold that allows SSR to occur. Due to the similarities between the SSR model and populations of neurons, and the observation that sensory neurons often operate in very low SNR conditions, we hypothesize that SSR may have a significant role to play in the encoding of signals by sensory neurons.

ACKNOWLEDGMENTS

This work was funded by the Australian Research Council, The Australian Academy of Science, the Cooperative Research Centre for Sensor Signal and Information Processing (Australia), and the Leverhulme Trust, grant number F/00 215/J and we thank them for their support.

REFERENCES

1. L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni, "Stochastic resonance," *Reviews of Modern Physics* **70**(1), pp. 223–287, 1998.
2. S. Kay, "Can detectability be improved by adding noise?" *IEEE Signal Processing Letters* **7**(1), pp. 8–10, 2000.
3. D. Abbott, "Overview: Unsolved problems of noise and fluctuations," *Chaos* **11**(3), pp. 526–538, 2001.
4. S. Fauve and F. Heslot, "Stochastic resonance in a bistable system," *Physics Letters A* **97A**(1-2), pp. 5–7, 1983.
5. B. McNamara, K. Wiesenfeld, and R. Roy, "Observation of stochastic resonance in a ring laser," *Physical Review Letters* **60**(25), pp. 2626–2629, 1988.
6. A. Longtin, "Stochastic resonance in neuron models," *Journal of Statistical Physics* **70**(1-2), pp. 309–327, 1993.

7. K. Wiesenfeld and F. Moss, "Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDS," *Nature* **373**(6509), pp. 33–36, 1995.
8. S. M. Bezrukov and I. Voydanoy, "Noise-induced enhancement of signal transduction across voltage-dependent ion channels," *Nature* **378**(6555), pp. 362–364, 1995.
9. Z. Gingl, L. B. Kiss, and F. Moss, "Non-dynamical stochastic resonance: Theory and experiments with white and arbitrarily coloured noise," *EuroPhysics Letters* **29**(3), pp. 191–196, 1995.
10. L. Gammaitoni, "Stochastic resonance and the dithering effect in threshold physical systems," *Physical Review E* **52**(5), pp. 4691–4698, 1995.
11. A. Bulsara and A. Zador, "Threshold detection of wideband signals: A noise induced maximum in the mutual information," *Physical Review E* **54**(3), pp. R2185–R2188, 1996.
12. N. G. Stocks, "Suprathreshold stochastic resonance in multilevel threshold systems," *Physical Review Letters* **84**(11), pp. 2310–2313, 2000.
13. N. G. Stocks, "Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance," *Physical Review E* **63**(4), 041114, 2001.
14. R. A. Wannamaker, S. P. Lipshitz, and J. Vanderkooy, "Stochastic resonance as dithering," *Physical Review E* **61**(1), pp. 233–236, 2000.
15. N. G. Stocks and R. Mannella, "Generic noise enhanced coding in neuronal arrays," *Physical Review E* **64**, 030902(R), 2001.
16. R. P. Morse, D. Allingham, and N. G. Stocks, "An information-theoretic approach to cochlear implant coding," in *UPON 2002: Third International Conference on Unsolved Problems of Noise and Fluctuations in Physics, Biology, and High Technology*, S. M. Bezrukov, ed., *AIP Conference Proceedings* **665**, pp. 125–132, 2002.
17. N. G. Stocks, D. Allingham, and R. P. Morse, "The application of suprathreshold stochastic resonance to cochlear implant coding," *Fluctuation and Noise Letters* **2**(3), pp. L169–L181, 2002.
18. D. Allingham, N. G. Stocks, R. P. Morse, and G. F. Meyer, "Noise enhanced information transmission in a model of multichannel cochlear implantation," in *Proc. SPIE Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, D. Abbott, S. M. Bezrukov, A. Der, and A. Sánchez, eds., **5467**, pp. 139–148, 2004.
19. M. D. McDonnell, N. G. Stocks, C. E. M. Pearce, and D. Abbott, "Point singularities and suprathreshold stochastic resonance in optimal coding." Preprint arXiv:cond-mat/0409528, 2004.
20. M. DeWeese and W. Bialek, "Information flow in sensory neurons," *Il Nuovo Cimento* **17**(7-8), pp. 733–741, 1995.
21. R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory* **44**(6), pp. 2325–2383, 1998.
22. A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publishers, 1992.
23. T. Berger and J. D. Gibson, "Lossy source coding," *IEEE Transactions on Information Theory* **44**(6), pp. 2693–2723, 1998.
24. M. D. McDonnell, D. Abbott, and C. E. M. Pearce, "An analysis of noise enhanced information transmission in an array of comparators," *Microelectronics Journal* **33**(12), pp. 1079–1089, 2002.
25. M. D. McDonnell, N. G. Stocks, C. Pearce, and D. Abbott, "Analog to digital conversion using suprathreshold stochastic resonance," in *Proc. SPIE BioMEMS and Nanoengineering*, In Press, 2005.
26. J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1999.