ASYMPTOTIC REFINEMENTS IN BAYESIAN DISTRIBUTED DETECTION[1]

Adrian Papamarcou and Po-Ning Chen

Electrical Engineering Department and Institute for Systems Research
University of Maryland, College Park MD 20742

Abstract

The performance of a parallel distributed detection system is investigated as the number of sensors tends to infinity. It is assumed that the i.i.d. sensor data are quantized locally into m-ary messages and transmitted to the fusion center for Bayesian binary hypothesis testing. Large deviations techniques are employed to show that the equivalence of absolutely optimal and best identical-quantizer systems is not limited to error exponents, but extends to the actual Bayes error probabilities up to a multiplicative constant. This is true as long as the two hypotheses are mutually absolutely continuous; no further assumptions, such as boundedness of second moments of the post-quantization log-likelihood ratio, are needed.

Summary

Consider a parallel distributed system consisting of n geographically dispersed sensors, noiseless one-way communication links, and a fusion center. Each sensor makes an observation (denoted by Y_i) of a random source, quantizes Y_i into an m-ary message $U_i = g_i(Y_i)$, and then transmits U_i to the fusion center. Upon receipt of $(U_1, \ldots, U_n)$, the fusion center performs a binary hypothesis test (H_0 against H_1) about the nature of the random source. A Bayesian setup is assumed throughout, and the Bayes error probability is denoted by $\gamma_n(\pi)$, where π is the prior probability of H_0.

It was shown by Tsitsiklis [1] that even when the observations are i.i.d., the optimal m-ary quantizers g_i need not be identical. Thus the absolutely optimal system (*) does not, in general, coincide with the best identical-quantizer system (°). Since the latter is much easier to design than the former, it is natural to seek an estimate of the performance loss resulting from using identical quantizers.

Tsitsiklis supplied a result of this type in the i.i.d. case by showing, under a fairly general assumption, that the two systems are asymptotically exponentially equivalent. More precisely, if P and Q are mutually absolutely continuous distributions of the i.i.d. observations under H_0 and H_1 respectively, then

$$\lim_{n\to\infty} -\frac{1}{n}\log\gamma_n^*(\pi) \;=\; \lim_{n\to\infty} -\frac{1}{n}\log\gamma_n^\circ(\pi)$$

(i.e., the two error exponents coincide) provided that the second moments—under P and Q—of the post-quantization log-likelihood ratio $\log[P_g(U)/Q_g(U)]$ are bounded as the quantizer mapping g varies. The optimal error exponent is the supremum (over g) of the Chernoff exponent associated with the m-ary post-quantization distributions P_g and Q_g. It has also been shown that this supremum is achieved by a g^* taken from the class of deterministic likelihood-ratio quantizers; and that such quantizers are optimal in the nonasymptotic (fixed n) setting.

Two questions arose from Tsitsiklis' work:

1. Is the aforementioned boundedness assumption really necessary?
2. Does the nonnegative quantity $\log[\gamma_n^\circ(\pi)/\gamma_n^*(\pi)]$ admit an upper bound tighter than $O(n)$ (implied by the equality of error exponents)?

Tsitsiklis [2] conjectured that the answer to the first question is negative, and in this paper we give proof to his conjecture. As to the second question, we show that the upper bound $O(n)$ on $\log[\gamma_n^\circ(\pi)/\gamma_n^*(\pi)]$ can be tightened to $O(1)$, hence the ratio $\gamma_n^\circ(\pi)/\gamma_n^*(\pi)$ is bounded from above (trivially, it is also lower-bounded by unity). We therefore have:

Theorem *If P and Q are mutually absolutely continuous, then for all $\pi \in (0,1)$,*

$$\limsup_{n\to\infty} \frac{\gamma_n^\circ(\pi)}{\gamma_n^*(\pi)} \;<\; \infty\,. \qquad \square$$

We employ large deviations techniques for proving this theorem. Using a refinement (due to Esseen) of the central limit theorem for independent but not identically distributed summands, we show the following: if all quantizers in the optimal system are "regular," in that they yield—at their output—log-likelihood ratios that satisfy certain uniform boundedness constraints, then the Bayes error probability $\gamma_n^*(\pi)$ can be lower-bounded by $cn^{-1/2}\exp\{-\rho_m n\}$, where ρ_m is the optimal error exponent. The same expression—only with a larger value of c—is also an upper bound on $\gamma_n^\circ(\pi)$, so the conclusion that $\gamma_n^\circ(\pi)/\gamma_n^*(\pi)$ is bounded from above is close at hand. It remains to show that the number of "irregular" quantizers in the optimal system is bounded. As it turns out, these quantizers yield an error exponent smaller than ρ_m, and thus heuristically, they can only exist in small numbers. We give rigorous proof to this fact using a technical argument based on conditioning.

Our simulations of Bayesian distributed detection have shown that the ratio $\gamma_n^\circ(\pi)/\gamma_n^*(\pi)$ is in many instances close to unity. It is quite possible that under conditions as yet unknown to us, the ratio $\gamma_n^\circ(\pi)/\gamma_n^*(\pi)$ tends to unity as n approaches infinity. This, however, is not true in general, and we give a counterexample in which the ratio $\gamma_n^\circ(\pi)/\gamma_n^*(\pi)$ is greater than $r > 1$ infinitely often in n.

References

[1] J. N. Tsitsiklis. Decentralized detection by a large number of sensors. *Mathematics of Control, Signals and Systems,* 1(2): 167–182, 1988.

[2] J. N. Tsitsiklis. On threshold rules in decentralized detection. In *Proc. 25th Conference on Decision and Control,* pages 232–236, Athens, Greece, 1986.

[1]Research supported by the Institute for Systems Research (a National Science Foundation Engineering Research Center) at the University of Maryland, College Park.