CHI-SQUARE ORACLE INEQUALITIES

IAIN M. JOHNSTONE[1]

Stanford University

We study soft threshold estimates of the non-centrality parameter ξ of a non-central $\chi_d^2(\xi)$ distribution, of interest, for example, in estimation of the squared length of the mean of a Gaussian vector. Mean squared error and oracle bounds, both upper and lower, are derived for *all* degrees of freedom d. These bounds are remarkably similar to those in the limiting Gaussian shift case. In nonparametric estimation of $\int f^2$, a dyadic block implementation of these ideas leads to an alternate proof of the optimal adaptivity result of Efromovich and Low.

AMS subject classifications: 62E17, 62F11, 62G07, 62G05.
Keywords and phrases: Quadratic Functionals, Adaptive Estimation, Gaussian sequence model, Efficient estimation, non-central chi-square.

1 Introduction

The aim of this paper is to develop thresholding tools for estimation of certain quadratic functionals. We begin in a finite dimensional setting, with the estimation of the squared length of the mean of a Gaussian vector with spherical covariance. The transition from linear to quadratic functionals of the data entails a shift from Gaussian to (non-central) chi-squared distributions $\chi_d^2(\xi)$ and it is the non-centrality parameter ξ that we now seek to estimate. It turns out that (soft) threshold estimators of the noncentrality parameter have mean squared error properties which, after appropriate scaling, very closely match those of the Gaussian shift model. This might be expected for large d, but this is not solely an asymptotic phenomenon – the detailed structure of the chi-squared distribution family allows relatively sharp bounds to be established for the full range of degrees of freedom d.

We develop oracle inequalities which show that thresholding of the natural unbiased estimator of ξ at $\sqrt{2 \log d}$ standard deviations (according to *central* χ_d^2) leads to an estimator of the non-centrality parameter that is within a multiplicative factor $2 \log d + \epsilon_d$ of an 'ideal' estimator that can use knowledge of ξ to choose between an unbiased rule or simply estimating zero. These results are outlined in Section 2.

Section 3 shows that the multiplicative $2 \log d$ penalty is sharp for large degrees of freedom d, essentially by reduction to a limiting Gaussian shift problem.

[1]This research was supported by NSF DMS 9505151 and ANU.