Available online at www.sciencedirect.com
ScienceDirect
Journal of Multivariate Analysis 99 (2008) 1362–1382
Journal of Multivariate Analysis
www.elsevier.com/locate/jmva

The centred parametrization for the multivariate skew-normal distribution

Reinaldo B. Arellano-Valle[a], Adelchi Azzalini[b,*]

[a] *Departamento de Estadística, Pontificia Universidad Católica de Chile, Santiago, Chile*
[b] *Dipartimento di Scienze Statistiche, Università di Padova, Italy*

Available online 12 February 2008

Abstract

For statistical inference connected to the scalar skew-normal distribution, it is known that the so-called centred parametrization provides a more convenient parametrization than the one commonly employed for writing the density function. We extend the definition of the centred parametrization to the multivariate case, and study the corresponding information matrix.

AMS 2000 subject classifications: 62F12; 62H12; 62E20

Keywords: Index of skewness; Skew-normal distribution; Parametrization; Singular information matrix

1. Background and motivation

1.1. The skew-normal distribution

In recent years, there has been increasing interest in a scheme for the genesis of families of distributions whose main features are collected in the book edited by Genton [13] and in the review paper by Azzalini [5]. The more emblematic and in a sense simplest representative of these distributions is the so-called skew-normal (SN) distribution, whose density function is, in the one-dimensional case,

$$f_1(x;\xi,\omega^2,\alpha) = 2\,\omega^{-1}\,\phi\left(\frac{x-\xi}{\omega}\right)\,\Phi\left\{\alpha\left(\frac{x-\xi}{\omega}\right)\right\},\quad (x\in\mathbb{R}), \tag{1}$$

* Corresponding author.
E-mail address: azzalini@stat.unipd.it (A. Azzalini).

0047-259X/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmva.2008.01.020

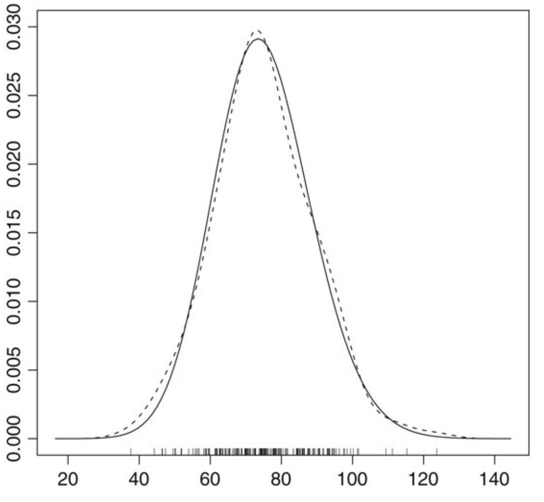

Fig. 1. AIS weight data: the ticks on the horizontal axis correspond to the individual data values, the dashed curve represents the nonparametric estimate of the density function, the continuous curve is the SN density fitted by maximum likelihood.

where ϕ and Φ denote the $N(0, 1)$ density and distribution function, respectively, and ξ, ω and α are location, scale and shape parameters, respectively ($\xi, \alpha \in \mathbb{R}$, $\omega \in \mathbb{R}^+$). The SN family forms a superset of the normal family, which corresponds to the choice $\alpha = 0$; with other values of α, a skewed density is obtained. Various other formal connections between (1) and the normal distribution are recalled in the above-mentioned references.

From the above-quoted general sources, it emerges that substantially more work has been oriented towards the probabilistic aspects of this approach and appreciably less to the statistical ones. Part of the explanation of this fact lies in the striking simplicity of treatment offered by this approach on the mathematical and probabilistic side, while the corresponding statistical work is, by contrast, surprisingly somewhat problematic.

Among these problematic aspects, a peculiar one concerns the behaviour of the likelihood function and other related quantities for a sample from the SN distribution in the neighbourhood of $\alpha = 0$, which is a value of particular relevance since there the SN family reduces to the normal one. We anticipate that the term 'neighbourhood' must be interpreted in a fairly 'wide' sense, as we shall see shortly.

Consider for illustration the AIS (Australian Institute of Sport) data, repeatedly employed in this stream of literature. These data contain a number of biomedical measurements on 202 Australian athletes, of which we shall use here only the weight (in kg). Fig. 1 reports on the horizontal axis the individual data values, and it displays a nonparametric estimate of the density function obtained by the kernel method, as well as the parametric curve of type (1) selected by maximum likelihood estimation (MLE), such that $(\hat{\xi}, \hat{\omega}^2, \hat{\alpha}) = (64.07, 17.68^2, 1.23)$. The general impression perceived from Fig. 1 is that the SN curve follows quite closely the nonparametric estimate of the density, and both curves indicate a mild departure from normality because of some skewness to the right.

While the fitting process is satisfactory in the sense of an overall indication of adequate parametric fit, there are other aspects which are less pleasant, due to some anomalies of the log-likelihood function. For a random sample $y_1, \ldots, y_n$ from distribution (1), the log-likelihood