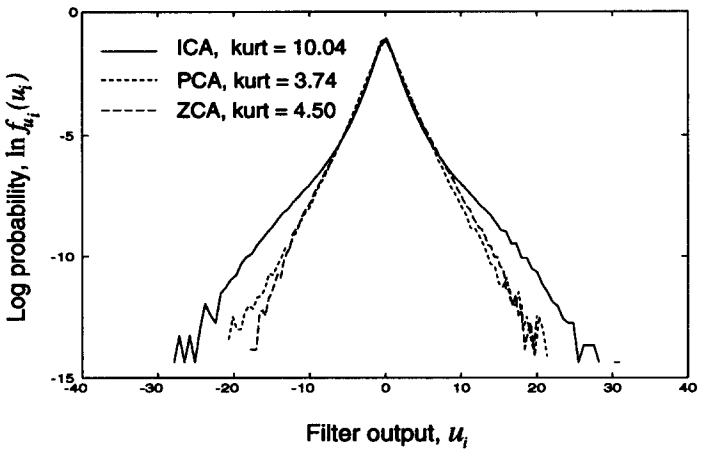

FIGURE 5. Log distributions of univariate statistics of the outputs of ICA, ZCA and PCA filters, averaged over all filters of each type. All three are approximately double-exponential distributions, but the more kurtotic ICA distribution is slightly peakier and has a longer tail, showing that it is *sparser* than the others. This distribution (and the 2-D ones in Fig. 6), although averaged over the outputs of all filters, are extremely similar to the distributions output by individual filters (respectively, pairs of filters). The only exception is the DC-filter (top left in Fig. 4) which has a more gaussian distribution.

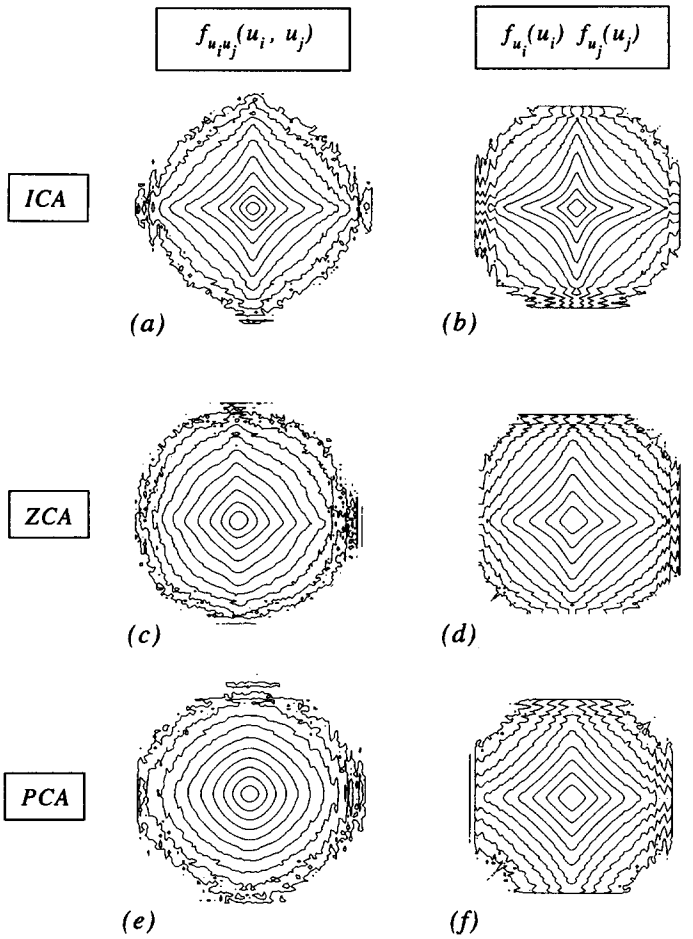

FIGURE 6. Contour plots of log distributions of pairwise statistics of the outputs of ICA, ZCA and PCA filters. Left column: joint log distributions averaged over all pairs of output filters of each type, and all images. Right column: product of marginal (univariate) distributions. The ICA solution best satisfies the independence criterion that the joint distribution has the same form as the product of the marginal distributions.

*The definition of "localized" causes some ambiguity here. While our ICA basis functions contain non-zero values all over the domain of the filter, their "contrast energy" occurs along one oriented local patch. PCA filters, on the other hand, are "more non-local" since neither of these conditions are satisfied.

and the basis functions in Olshausen and Field's Fig. 4. The ICA basis functions in Fig. 3(e) are oriented, but not localized and therefore it is difficult to observe any multiscale properties.* However, when we ran the ICA algorithm on Olshausen's images, which were prepro-