Using Principal Component Analysis to Fit PSF Variation in Large Fields

Abstract

High precision cosmology with weak lensing requires significant improvements in the measurement of galaxy shapes. We present results from a new analysis technique that reduces the contribution of systematic errors in shape measurements. This technique addresses the difficulty in modeling the PSF variation given the relatively small number of stars per image. With of order 100 stars per image, many of which are quite noisy, we have generally been limited to about 4th order polynomial fits. As the PSF often varies more quickly than this in some regions of wide-field images, the resulting errors in the fits have caused B-mode systematics in the lensing analyses. Performing a principal component analysis of the variation allows us to use stars from many images at once to accurately fit a higher order function for the variation. This technique has significantly reduced our B-mode contamination for our CTIO lensing survey. While the problem is not nearly as severe in space, the precision demanded by future lensing surveys will likely require a technique such as this, which is easily applicable to space-based telescopes.

Results

The best test for residual systematics in weak lensing surveys so far has been to look for a B-mode component to the shear pattern. Weak lensing should only produce a curl-free E-mode component, so any B-mode represents a systematic error of some sort. Every weak lensing study which has tested for B-mode has detected some, although usually at a level somewhat smaller than the E-mode signal. For our CTIO survey data, applying the PCA technique has reduced the detected B-mode at all scales above 2 or 3 arcmin.

Below we show the results for a particular measurement known as the aperture mass statistic. The **blue points are the E-mode signal**, and the **red points are the B-mode contamination**. Points separated by at least one other point are uncorrelated. The E-mode error bars include the contamination from the B-mod. The **black curves are the prediction for WMap's best-fit ΛCDM cosmology**.

The left plot shows the analysis without our PCA technique, and shows B-mode contamination at scales less than 10 arcmin. The right plot shows the new analysis including this technique. **Almost all the B-mode measurements are now consistent with zero.**

Without PCA

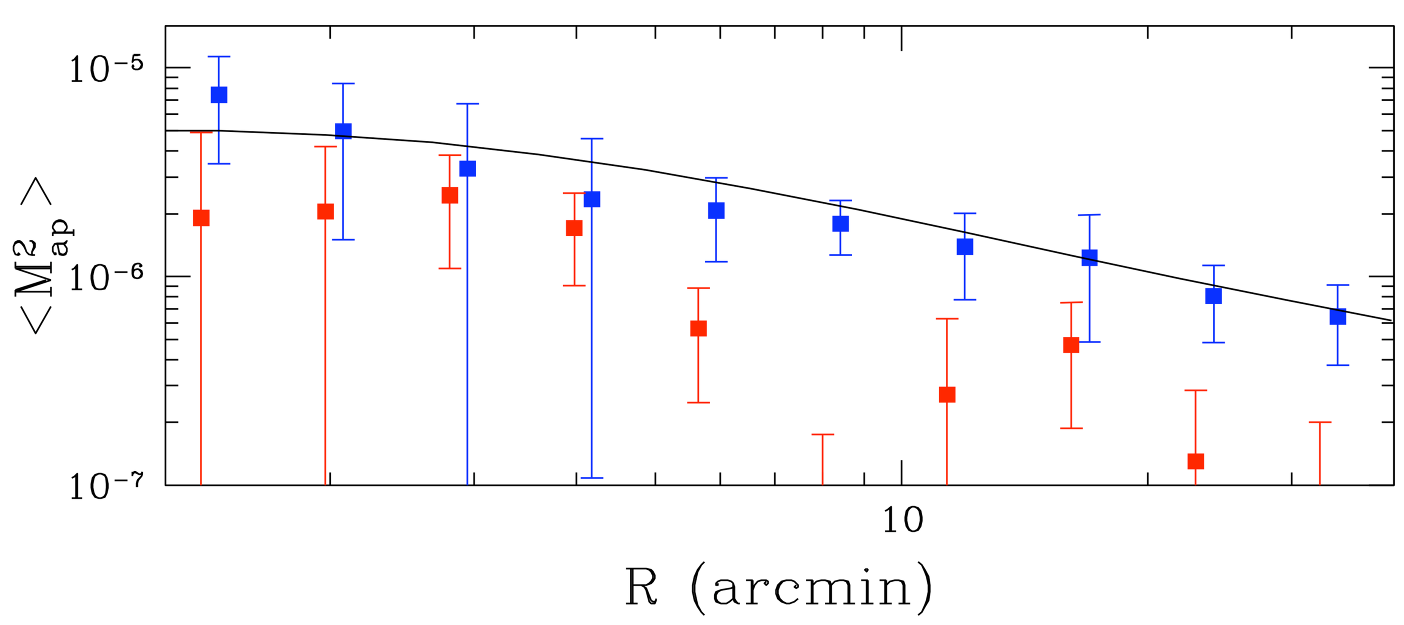

With PCA

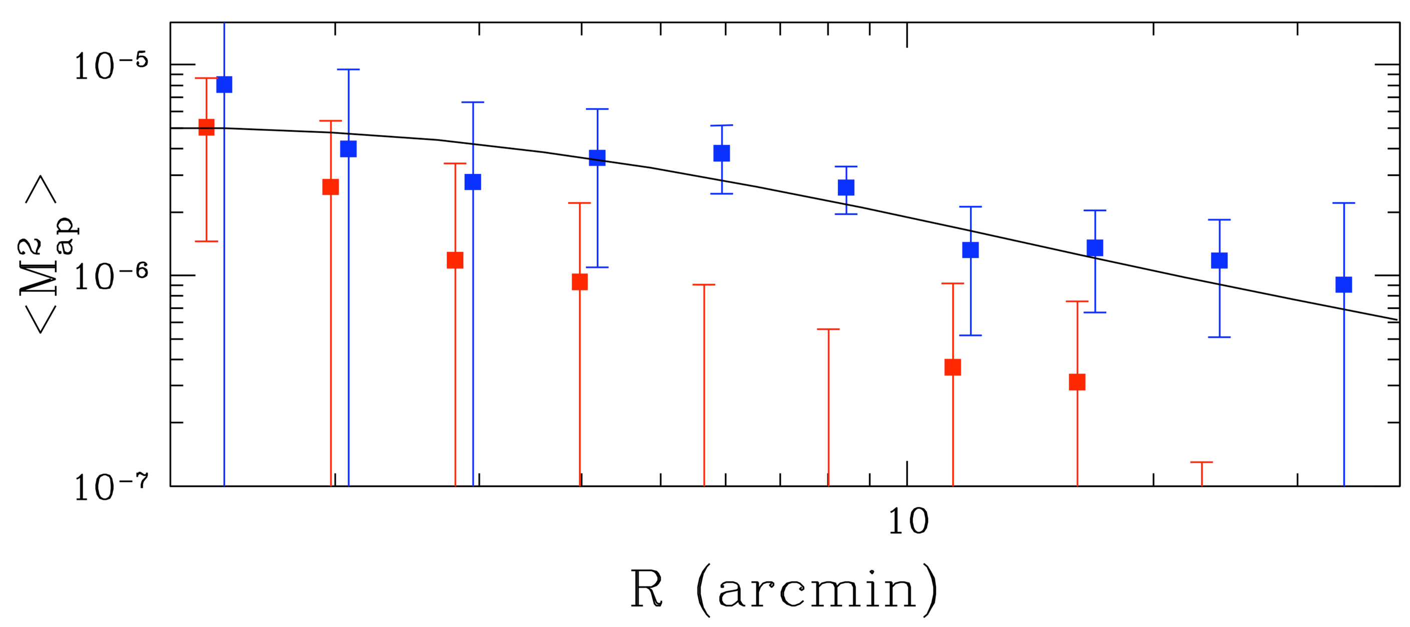

The Method in Detail

- First, we fit a low order polynomial (in x and y) to the PSF shape for each image. The exact details of this step are not too important. Basically, we just want a set of numbers which describes the PSF pattern.

- Arrange these numbers into a matrix where each row of the matrix corresponds to a different exposure, and the elements in the row are the numbers from the above fit.

- Perform a singular value decomposition of this matrix:
  $$M = USV$$
  also known as principal component analysis. (U and V are unitary matrices, S is diagonal.)

  The rows of V are called the principal components. Each row of M can be written as a weighted sum of these components:
  $$M(i,*) = \sum_k U(i,k)S(k,k)V(k,*)$$
  Most of the variation in the rows of M is described by the first principal component ($k = 1$). Thus, the value $U(i,0)$ is really telling us the focus position of that exposure. Call this value f_i:
  $$f_i = U(i,0)$$

- Next we fit the PSF shapes as a function of the chip positions x and y and the focus value for that exposure:
  $$e_{\mathrm{PSF}}(x,y,i) = F(x,y,f_i)$$
  The form of this function is somewhat arbitrary. We choose F to be polynomial in x and y, and linear in f_i:
  $$e_{\mathrm{PSF}}(x,y,i) = P_0^{(n)}(x,y) + f_i P_1^{(n)}(x,y)$$
  where P_0 and P_1 are n[th] order polynomial functions.

  We don't have to worry too much about higher order terms in f_i, since if the PSF varied as some function of f_i, instead of linearly, then the principal component decomposition would have picked those values out as the f_i instead.

- However, while the first principal component is indeed fairly dominant, the next several components are not insignificant. So the obvious generalization is to use the first several values of $U(i,k)$ for each i as a set of "focus" values f_{ik}:
  $$e_{\mathrm{PSF}}(x,y,i) = \sum_k f_{ik} P_k^{(n)}(x,y)$$
  We drop the 0 term (with no f_{ik} coefficient), because to the extent that it was important, it will be represented as a linear combination of some of the f_{ik} values.

- The real advantage to this fit is that we can now use all of the stars from every image to constrain the fit parameters. There are about 10 (or however many f_{ik} we choose to use) times as many coefficients to constrain, but we have hundreds of times as many stars to do the constraining. This lets us take n (the order of the polynomial fits) to be significantly higher than we could doing one image at a time.

- Also, this technique allows for better outlier rejection, since there is less worry about rejecting all of the stars in a region where the PSF changes quickly.

- Finally, once we have fit all the coefficients, we convert back to a simple polynomial for each image by summing over the f_{ik} coefficients.
  $$e_{\mathrm{PSF},i}(x,y) = \widetilde{P}_i^{(n)}(x,y) = \sum_k f_{ik} P_k^{(n)}(x,y)$$
  This function $\widetilde{P}_i^{(n)}(x,y)$ is then used for all subsequent analysis.

- In space, the principal variation in the PSF may not be due to focus errors, since focusing is somewhat easier than for ground-based telescopes. However, there may be effects from the sun's position relative to the pointing, or the temperature of the mirror, or some other unforeseen factor. This technique automatically picks out the important effects and lets you correct for them.

Stellar Whisker Plots

At right are three "whisker plots" from our CTIO weak lensing survey. Each line ("whisker") represents the shape of an observed star. The line is oriented in the same direction as the PSF shape, and the length is proportional to its ellipticity.

The leftmost plot represents the most out of focus image in one direction, and the rightmost plot represents the extreme in the other direction. The middle plot is far more typical, but as you can see, the four chips are not precisely coplanar, so they do not all come into focus at exactly the same time.

Also evident in these plots is the noisiness of the shape measurements. This limits us to approximately 4th order fits (in each of x and y) with the typical 100 or so stars per image.

Finally, we note that observing a dense stellar field (as proposed by Hoekstra, 2004) is not sufficient for our survey, since the PSF pattern changes dramatically with the relative focus position. Using the high order fit from some "average" image will not help any image with a different focus value.

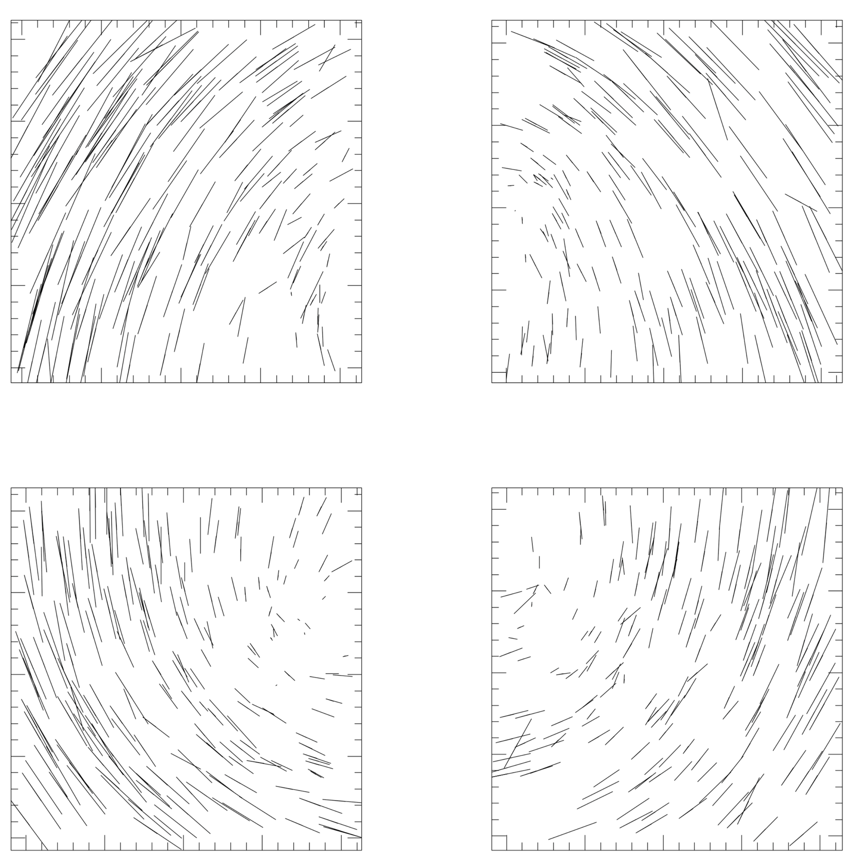

Focus too low

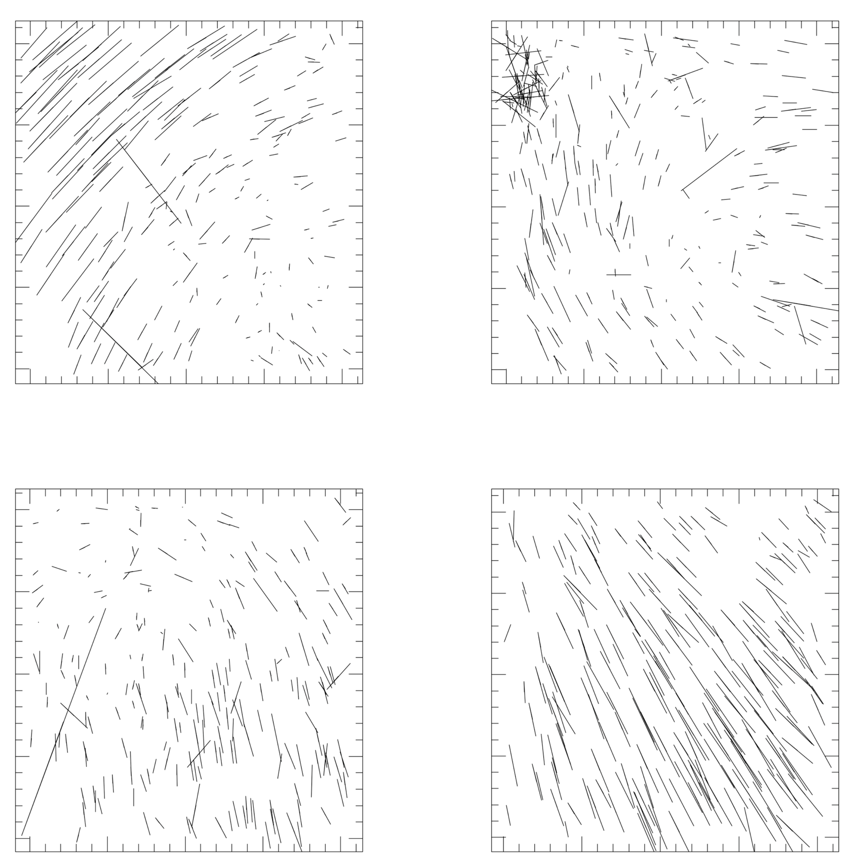

Focus (roughly) correct

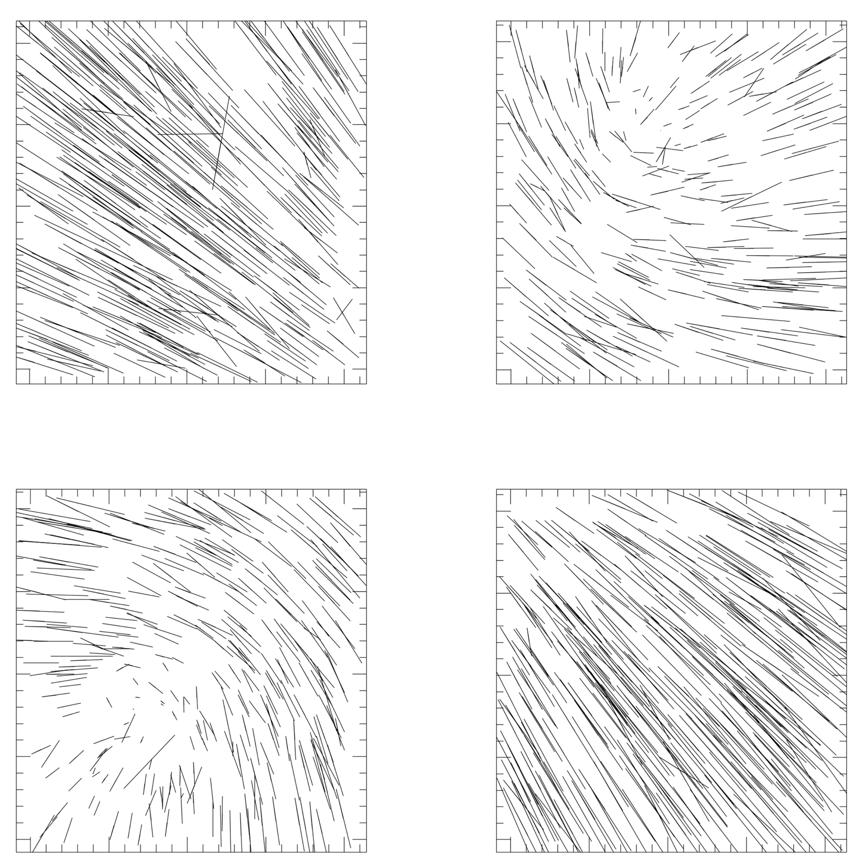

Focus too high