

Deconvolution with shapelets

P. Melchior, R. Andrae, M. Maturi, and M. Bartelmann

Zentrum für Astronomie, ITA, Universität Heidelberg, Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany
 e-mail: pmelchior@ita.uni-heidelberg.de

Received 27 June 2008 / Accepted 12 September 2008

ABSTRACT

Aims. We seek a shapelet-based scheme for deconvolving galaxy images from the PSF that leads to unbiased shear measurements.

Methods. Based on the analytic formulation of convolution in shapelet space, we constructed a procedure to recover the unconvolved shapelet coefficients under the assumption that the PSF is perfectly known. Using specific simulations, we test this approach and compare it to other published approaches.

Results. We show that convolution in shapelet space leads to a shapelet model of order $n_{\max}^h = n_{\max}^g + n_{\max}^f$, with n_{\max}^f and n_{\max}^g being the maximum orders of the intrinsic galaxy and the PSF models, respectively. Deconvolution is hence a transformation that maps a certain number of convolved coefficients onto a generally smaller number of deconvolved coefficients. By inferring the latter number from data, we construct the maximum-likelihood solution for this transformation and obtain unbiased shear estimates with a remarkable amount of noise reduction compared to established approaches. This finding is particularly valid for complicated PSF models and low S/N images, which renders our approach suitable for typical weak-lensing conditions.

Key words. techniques: image processing – gravitational lensing – methods: data analysis

1. Introduction

Shapelets have been proposed as an orthonormal set of two-dimensional functions to quantify shapes of galaxy images (Refregier 2003). They have several convenient mathematical properties that suggest their use in measurements of weak gravitational lensing (Refregier & Bacon 2003). Their application to data with low signal-to-noise ratio, however, is hampered by several problems. First, their enormous flexibility allows shapelets to represent random noise patterns well, which generates their tendency to escape into high orders and to overfit noisy images. Second, this same property affects not only the order, but also the spatial scale of the best-fitting shapelet model. The scale sizes of the shapelets decomposing noisy images are thus likely to be too high. These drawbacks are inherent properties of the shapelet method and have to be addressed to effectively use them in astronomical image processing. At first sight, it seems natural to limit the order of the shapelets from above to avoid overfitting and to limit the scale from above to prevent shapelets from creeping into the noise. On the other hand, convolution with the PSF alters the shape and spatial extent of imaged objects, leaving an imprint on the maximum shapelet order and the scale size.

Guided by these considerations, we address the question of how the maximum order and the spatial scale of a shapelet decomposition should be determined in cases where PSF convolution significantly modifies the object's appearance? This question aims at applications of shapelets in weak-lensing measurements, where the situation is particularly delicate because typically very small and noisy images are convolved with structured, incompletely known PSF kernels with scales similar to those of the imaged objects. How can deconvolution schemes be constructed in this case to find significant and unbiased shear estimates?

We show in Sect. 2 and in the Appendix that mathematical sum rules exist for the shapelet orders and the squared spatial scales of original image, kernel, and convolved image. We then proceed in Sect. 3 to devise an algorithm respecting these sum rules as closely as possible, which leads to a deconvolution scheme based on the (possibly weighed) pseudo-inverse of a rectangular rather than the inverse of a quadratic convolution matrix. In Sect. 4, we demonstrate by means of simulations with different signal-to-noise levels and PSF kernels that our algorithm does indeed perform very well and that it leads in most realistic cases to substantially improved results compared to previously proposed methods. Our conclusions are summarized in Sect. 5.

2. Convolution in shapelet space

A two-dimensional function $f(\mathbf{x})$, e.g. a galaxy image, is decomposed into a set of shapelet modes by projection,

$$f_n = \int_{-\infty}^{\infty} d^2x f(\mathbf{x}) B_n(\mathbf{x}; \beta_f), \quad (1)$$

where $\mathbf{n} = (n_1, n_2)$ is a two-dimensional index and β_f is called *scale size*. The two-dimensional shapelet basis function

$$B_n(\mathbf{x}; \beta_f) = \beta_f^{-1} \phi_{n_1}(\beta_f^{-1} x_1) \phi_{n_2}(\beta_f^{-1} x_2), \quad (2)$$

is related to the one-dimensional Gauss-Hermite polynomial

$$\phi_n(x) = [2^n \pi^{1/2} n!]^{-1/2} H_n(x) e^{-x^2/2}, \quad (3)$$

with $H_n(x)$ the Hermite polynomial of order n . From the coefficients one can then reconstruct a shapelet model

$$\tilde{f}(\mathbf{x}) = \sum_n^{n_{\max}} f_n B_n(\mathbf{x}; \beta_f). \quad (4)$$

The number of shapelet modes – often expressed in terms of the maximum shapelet order $n_{\max} = \max(n_1, n_2)$ – and the scale size have to be determined by an optimization algorithm (Massey & Refregier 2005; Melchior et al. 2007), which minimizes the modulus of the residuals $f - \tilde{f}$, or from empirical relations based on other measures of the object like FWHM or major and minor axes (Chang et al. 2004; Kuijken 2006).

A convolution

$$h(\mathbf{x}) \equiv (f \star g)(\mathbf{x}) \equiv \int_{-\infty}^{\infty} d^2 x' f(\mathbf{x}') g(\mathbf{x} - \mathbf{x}') \quad (5)$$

can be performed analytically in shapelet space. According to Eq. (1), the functions f , g , and h are represented by sets of shapelet states f_n , g_n , and h_n with scale sizes β_f , β_g , and β_h .

Refregier (2003) and Refregier & Bacon (2003) showed that the coefficients of the convolved image $h(\mathbf{x})$ are given by

$$h_n = \sum_{m,l} C_{n,m,l}(\beta_f, \beta_g, \beta_h) f_m g_l = \sum_m P_{n,m}(\beta_f, \beta_g, \beta_h) f_m. \quad (6)$$

Formally, this is an ordinary matrix equation,

$$h = P f, \quad (7)$$

hence $P_{n,m} \equiv \sum_l C_{n,m,l} g_l$ is called *convolution matrix*. The value of $C_{n,m,l}(\beta_f, \beta_g, \beta_h)$ and thus the entire convolution matrix can be computed analytically.

However, there is no clear statement of the scale size β_h and, in particular, of the maximum order n_{\max}^h of the convolved object h . In appendix A we prove that the so-called *natural choice* (Refregier 2003)

$$\beta_h^2 = \beta_f^2 + \beta_g^2 \quad (8)$$

is indeed the correct choice for β_h and that the maximum order of the convolved object is given by

$$n_{\max}^h = n_{\max}^f + n_{\max}^g. \quad (9)$$

While this result gives the highest possible mode of the convolved object that could contain power, it does not tell us whether it does indeed have power, as this depends primarily on the ratio of scales β_f/β_g entering $P_{n,m}$. This is demonstrated in Fig. 1, where we show the result of a convolution of a function that is given by a pure B_4 mode with a kernel represented by a pure B_2 mode. From this it becomes obvious that in a wide region around $\beta_f/\beta_g \simeq 1$, power is transferred to all even modes up to $n = 6$ (odd modes vanish because of parity, cf. Eq. (A.5) and the following discussion). If either $\beta_f \gg \beta_g$ or $\beta_g \gg \beta_f$, the highest order of the larger object is also the highest effective order of the convolved object. Thus, we can generalize Eq. (9),

$$n_{\max}^h = \begin{cases} n_{\max}^f (+1) & \beta_f \gg \beta_g \text{ (kernel negligible)} \\ n_{\max}^f + n_{\max}^g & \beta_f \simeq \beta_g \\ n_{\max}^g (+1) & \beta_f \ll \beta_g \text{ (kernel dominant)} \end{cases} \quad (10)$$

where the option (+1) is taken if required by parity.

For most cases, in particular for weak gravitational lensing, PSF and object scales are comparable, which means that we must not neglect the power transfer to higher modes.

3. How to deconvolve

While the correct values for β_h and n_{\max}^h are clarified now, it is not obvious how these results need to be used in deconvolving real data. We first comment on possible ways to undo the convolution and then discuss our finding in the light of measurement noise.




Fig. 1. One-dimensional convolution $h = B_4(x; \beta_f) \star B_2(x; \beta_g)$. We plot the modulus of h_n with n even (all odd modes have vanishing power), normalized by $\sum_n |h_n|$.

3.1. Deconvolution strategies

As Refregier & Bacon (2003) have already discussed, there are two ways to deconvolve from the PSF in shapelet space:

– inversion of the convolution matrix: according to Eq. (6), one can solve for the unconvolved coefficients;

$$f_m = \sum_n P_{m,n}^{-1} h_n; \quad (11)$$

– fit with the convolved basis system (Kuijken 1999; Massey & Refregier 2005): This modifies Eq. (4) such that it directly minimizes the residuals of $h(\mathbf{x})$ w.r.t. its shapelet model

$$\tilde{h}(\mathbf{x}) = \sum_n f_n \sum_m P_{n,m} B_m(\mathbf{x}; \beta_f). \quad (12)$$

The second method is generally applicable but slow because the convolution has to be applied at each iteration step of the decomposition process. On the other hand, the first approach reduces deconvolution to a single step after the shapelet decomposition and is therefore computationally more efficient.

According to Eq. (9), P is not square and thus not invertible as suggested by Eq. (11). To cope with this, we need to replace the inverse P^{-1} by the pseudo-inverse $P^\dagger \equiv (P^T P)^{-1} P^T$ such that the equation now reads as

$$f_m = \sum_n P_{m,n}^\dagger h_n. \quad (13)$$

What seems as a drawback at first glance is effectively beneficial. Conceptually, this is now the least-squares solution of Eq. (6), recovering the most probable unconvolved coefficients from the set of noisy convolved coefficients. The underlying assumption of Gaussian noise in the coefficients h_n holds for the most usual case of background-dominated images for which the pixel noise is Gaussian.

Correlations among the coefficients h_n , which arise from non-constant pixel weights or pixel correlations, can be accounted for by introducing the coefficient covariance matrix $\Sigma \equiv \langle (h_n - \langle h_n \rangle)(h_m - \langle h_m \rangle)^T \rangle$, which alters Eq. (7) to read

$$P^T \Sigma P f = P^T \Sigma h. \quad (14)$$

Maximizing the likelihood for recovering the correct unconvolved coefficients leads to the weighted pseudo-inverse

$$P_w^\dagger \equiv (P^T \Sigma P)^{-1} P^T \Sigma, \quad (15)$$

which replaces P^\dagger in Eq. (13).

However, both approaches (direct inversion, Eq. (11), or least-squares solution, Eq. (13)) would fail if P was rank-deficient. Refregier & Bacon (2003) argue that convolution with the PSF amounts to a projection of high-order modes onto low-order modes and therefore P can become singular. This is only true for very simple kernels (e.g. the Gaussian-shaped mode of order 0) with rather large scales. In fact, Eq. (10) tells us that convolution carries power from all available modes of f to modes up to order $n_{\max}^h \geq n_{\max}^f$, hence P is generally not rank-deficient. In practice we did not have problems in constructing P^{-1} or P^\dagger when using realistic kernels. We therefore see no hindrance in employing the matrix-inversion scheme and will use it in the course of this paper.

3.2. Measurement process and noise

Up to here, we have discussed (de-)convolution entirely in shapelet space, where this problem is now completely solved. For the following line of reasoning, we further assume that the kernel is perfectly known and can be described by a shapelet model.

Critical issues still arise at the transition from pixel to shapelet space. There are no intrinsic values of n_{\max}^f and β_f . Even if they existed, they would not be directly accessible to a measurement. While the first statement stems from trying to model a highly complicated galaxy or stellar shape with a potentially completely inappropriate function set, the second statement arises from pixelation and measurement noise occurring in the detector.

However, the pixelated version of the shape can be described by a shapelet model, with an accuracy that depends on the noise level and the pixel size. Consider for example a galaxy whose light distribution strictly follows a Sérsic profile. Modeling the cusp and the wide tails of this profile with the shapelet basis functions would require an infinite number of modes. But pixelation effectively removes the central singularity of the Sérsic profile and turns the continuous light distribution into a finite number of light measures, such that it is in principle describable by a finite number of shapelet coefficients. Pixel noise additionally limits the spatial region within which the tails of Sérsic profile remain noticeable, hence the number of required shapelet modes.

Consequently, shapelet implementations usually determine n_{\max} by some significance measure of the model (χ^2 in Massey & Refregier 2005; Melchior et al. 2007) or – similarly – fix n_{\max} at a value that seems reasonable for capturing the general features of the shape (e.g. Refregier & Bacon 2003; Kuijken 2006).

Figure 2 schematically highlights an important issue of a significance-based ansatz. When the power in a shapelet coefficient is lower than the power of the noise, it is considered insignificant, and the shapelet series is truncated at this mode (in Fig. 2, f_n may be limited to $n \leq 2$ and h_n to $n \leq 3$). Since convolution with a flux-normalized kernel does not change the overall flux or – as the shapelet decomposition is linear – the total coefficient power but generally increases the number of modes, the signal-to-noise ratio (S/N) of each individual coefficient is lowered on average. Thus, after convolution, more coefficients will be considered insignificant and will be disregarded.

This is equivalent to the action of a convolution in pixel space, where some objects' flux is distributed over a larger area. If the noise is independent of the convolution, demanding a




Fig. 2. Sketch of the effect of a convolution on the power of shapelet coefficients. The detailed shape of the curves is neither realistic nor important, but typical shapelet models show decreasing coefficient power with increasing order n . The noise regime (represented by the gray area) is constant in the case of uncorrelated noise.

certain S/N threshold results in a smaller number of significant pixels.

The main point here is that we try to measure h_n from data and from this f_n by employing Eq. (13). But if we truncate h_n too early, at an order $n_{\max}^h \ll n_{\max}^f + n_{\max}^g$, the resulting unconvolved coefficients f_n are expected to be biased even if the convolution kernel is perfectly described. The reason for this is that, by truncating, we assume that any higher-order coefficient is zero on average, while in reality it is non-zero, but just smaller than the noise limit. Every estimator formed from these coefficients is thus likely biased itself.

In turn, if we knew n_{\max}^f , we could go to the order demanded by Eq. (9) and the deconvolution would map many noise-dominated high-order coefficients back onto lower-order coefficients. This way, we would not cut off coefficient power, and our coefficient set would remain unbiased. Unfortunately, this approach comes at a price. First, the resulting shapelet models are often massively overfitted, and second, obtaining unbiased f_n requires the knowledge of n_{\max}^f . The first problem can be addressed by averaging over a sufficient number of galaxies, while the second one can indeed be achieved by checking the S/N of the recovered f_n after deconvolution. The average number of significant deconvolved coefficients gives an indication of the typical complexity of the imaged objects as they would be seen in a measurement without convolution, but using the particular detector characterized by its pixel size and noise level.

3.3. Unbiased deconvolution method

The previous consideration guides us to set up a deconvolution procedure that yields unbiased deconvolved coefficients. Again, we assume that the kernel can be perfectly modeled by shapelet series with finite order n_{\max}^g ¹.

- Given the noise level and the pixel size of the images, we initially guess \bar{n}_{\max}^f for each individual galaxy;
- we set the lower bounds $n_{\max}^h \geq n_{\max}^g + \bar{n}_{\max}^f$ and $\beta_h \geq \beta_g$;

¹ In reality, this assumption does not hold, in particular due to under-sampling of the PSF and pixel noise in those images from which the kernel model has to be inferred. Since g enters P (Eq. (7)), and P has to be inverted explicitly (Eq. (11)) or implicitly (Eq. (13)), a modest error in the kernel modeling can have severe implications for the quality of the deconvolution results. But this is also true for any other deconvolution method.




Fig. 3. Example of simulated galaxies used in our testbed: **a)** intrinsic galaxy model with $\beta_f = 2$ and flux equaling unity; **b)** after applying a shear $\gamma = (0.1, 0)$; **c)** after convolving with PSFb from Fig. 4 with $\beta_g = 2$. The bottom panels show **c)** after addition of Gaussian noise of zero mean and variance σ_n^2 : **d)** moderate noise, $\sigma_n = 10^{-4}$; **e)** high noise, $\sigma_n = 10^{-3}$; **f)** is the shapelet reconstruction of **e)**. The scale is logarithmic.

- we decompose each galaxy by minimizing the decomposition χ^2 under these constraints. A value of $n_{\max}^h > n_{\max}^g + \tilde{n}_{\max}^f$ is used only if $\chi^2 > 1$ otherwise. This yields h_n and β_h ;
- by inverting Eq. (8), we obtain $\tilde{\beta}_f$;
- using the maximum orders and scale sizes for f , g , and h in addition to g_n , we can form the convolution matrix P according to Eq. (7);
- by forming $P_{(w)}^{\dagger}$ and applying Eq. (13), we reconstruct \tilde{f}_n ;
- by propagating the coefficient errors from the decomposition through the same set of steps, we investigate the number of significant coefficients and should find \tilde{n}_{\max}^f if our initial guess was correct.

Given the demanded accuracy, it might be necessary to adjust the guess \tilde{n}_{\max}^f and reiterate the steps above. For the initial guess, it is inevitable to split the data set in magnitude bins, because the best value for \tilde{n}_{\max}^f clearly depends on the intrinsic brightness. Further splitting (according to apparent size or brightness profile etc.) may be advantageous, too.

4. Deconvolution microbenchmark

There exists a growing number of shapelet-based decomposition and deconvolution approaches published in the literature. In this section we show that the method proposed here is indeed capable of inferring unbiased unconvolved coefficients. Moreover, employing the least-squares solution given by Eq. (13) results in a considerable noise reduction, which is to be expected from this ansatz.

At first, we want to emphasize that the simulations we use in this section are highly simplistic. Their only purpose is to investigate how well a certain decomposition/deconvolution scheme can recover the unconvolved coefficients. By understanding the performance of different approaches, we acquire the knowledge for treating more realistic cases.

4.1. The testbed

The construction of simulated galaxy images is visualized in Fig. 3. As intrinsic function we use a polar shapelet model




Fig. 4. The kernels used in our benchmark: **a)** model of PSF2 from STEP1 (Heymans et al. 2006) with $n_{\max}^g = 4$; **b)** model of PSF3 from STEP1 with $n_{\max}^g = 4$; **c)** airy disk model with $n_{\max}^g = 6$; **d)** model from a raytracing simulation of a space-borne telescope’s PSF with $n_{\max}^g = 8$ and $n_{\max}^g = 12$ (shown here). The scale is logarithmic.

with $f_{0,0} = f_{2,0} = c$, where c is chosen such that the model has unit flux. Then, β_f is varied between 1.5 and 4. Given its ring-shaped appearance, this model is not overly realistic but also not too simple, and circularly symmetric. We apply a mild shear of $\gamma = (0.1, 0)$, thus populate coefficients of order $n_{\max}^f \leq 4$, and convolve with five different realistic kernels g (cf. Fig. 4) in shapelet space (employing Eqs. (8) and (9) with $1.5 \leq \beta_g \leq 6$). The pixelated version of the convolved object is then subject to N realizations of Gaussian noise with constant variance.

Each of these simulated galaxy images is decomposed into shapelets again, yielding h_n , using the code by Melchior et al. (2007), where the optimization is constrained by fixing either n_{\max}^h or β_h , or both. The observed coefficients h_n are then deconvolved from the kernel g .

As a diagnostic for the correctness of the deconvolved coefficients, we estimate the gravitational shear from the quadrupole moments Q_{ij} of the light distribution (Bartelmann & Schneider 2001),

$$\tilde{\gamma}_Q = \frac{1}{2} \mathcal{X} \equiv \frac{1}{2} \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22}}, \quad (16)$$

where Q_{ij} is computed as a linear combination of all available deconvolved coefficients (Bergé 2005; Melchior et al. 2007).

We investigate five different approaches that differ in the choice of n_{\max}^h , \tilde{n}_{\max}^f or the reconstruction of β_f . The different choices are summarized in Table 1.

FULL is the method we propose here (cf. Sect. 3.3), and for the following tests we set $\tilde{n}_{\max}^f = n_{\max}^f = 4$. SIGNIFIC is a variant of FULL, which bounds the decomposition order by the kernel order because coefficients beyond that are often insignificant, but makes use of our guess on \tilde{n}_{\max}^f .

SAME is similar to the one used by Kuijken (2006) with two differences. As discussed above, we employ the matrix inversion scheme (Eq. (11) since P is square for this method) instead of fitting the convolved shapelet basis functions, and in our implementation χ^2 is minimized with respect to a continuous parameter β_h , while Kuijken (2006) finds the best-fitting $\beta_h = 2^{n/8} \beta_g$ with some integer n . Without knowing the increase of shapelet orders due to convolution given by Eq. (9), this represents the best-defined deconvolution approach.

Refregier & Bacon (2003) state that the approach CONSTSCALE delivers the best results in their analysis. NMAX2, however, is an approach inspired by the naïve assumption that such a decomposition scheme catches the essential shear information without being affected by overfitting.

4.2. Performance with moderate noise

The first set of simulations comprises galaxy models with peak S/N between 45 and 220 with a median of ≈ 90 (an example

Table 1. Overview of the parameter choices of the investigated methods.

Name	n_{\max}^h ^a	$\tilde{\beta}_f$ ^b	\tilde{n}_{\max}^f ^c
FULL	$\geq n_{\max}^g + \tilde{n}_{\max}^f$	$\sqrt{\beta_h^2 - \beta_g^2}$	\tilde{n}_{\max}^f
SIGNIFIC	$\geq n_{\max}^g$	$\sqrt{\beta_h^2 - \beta_g^2}$	\tilde{n}_{\max}^f
SAME	n_{\max}^g	$\sqrt{\beta_h^2 - \beta_g^2}$	n_{\max}^g
CONSTSCALE	n_{\max}^g	β_h	n_{\max}^g
NMAX2	2	$\sqrt{\beta_h^2 - \beta_g^2}$	2

^a Order of the decomposed object; ^b estimate on the intrinsic scale of f ;
^c estimate on the intrinsic order of f .




Fig. 5. Recovered shear $\tilde{\gamma}$ and intrinsic scale size $\tilde{\beta}_f$ in a moderate-noise simulation ($\sigma_n = 10^{-4}$, $\beta_f = 3$, PSFb) as functions of the kernel scale β_g . In each panel the horizontal dashed line shows the true value of the quantity and the vertical dotted line shows the true value for β_f as reference. Errorbars (which are often too small to be visible) exhibit the standard deviation of the mean of $N = 100$ realizations. For visualization purposes, each method is slightly offset along β_g . Simulations with different PSF models or β_f are qualitatively equivalent.

is shown in Fig. 3d). For each value of β_f and β_g , we created $N = 100$ noise realizations. These high S/N values are more typical of galaxy morphology studies than of weak lensing, but we can see the effect of the convolution best. In this regime, problems with the deconvolution method immediately become apparent.

Considering Fig. 5, we can ascertain that FULL, SIGNIFIC, and SAME perform quite well, while CONSTSCALE and NMAX2

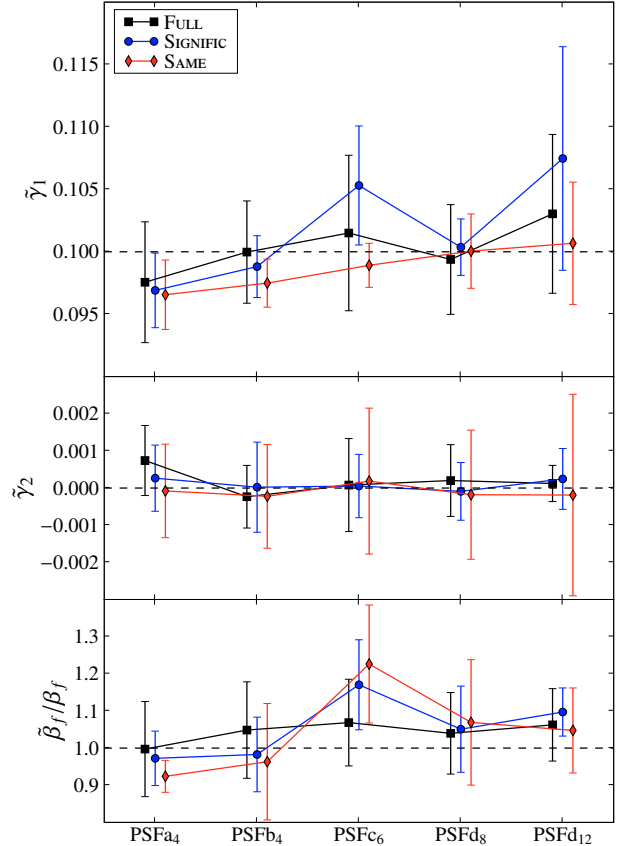


Fig. 6. Recovered shear $\tilde{\gamma}$ and intrinsic scale size $\tilde{\beta}_f$ (in units of the true scale size β_f) in the moderate-noise simulations in dependence on the PSF models from Fig. 4 (subscripts denote n_{\max}^g). Each data point represents the mean of the quantity for all available values of β_f and β_g (in total 60 independent combinations), and errorbars show the standard deviation of the mean. For visualization purposes, each method is slightly offset horizontally with respect to the others.

are clearly in trouble. This is not too surprising. By construction, NMAX2 truncates the shapelet series at $n_{\max}^h = 2$, hence misses all information contained in higher-order coefficients. One has to recall that the sheared model already has $n_{\max}^f = 4$, and after convolution with PSFb ($n_{\max}^g = 4$), it arrives at $n_{\max}^h = 8$. NMAX2 tries to undo the deconvolution with less information than contained in both sheared model and kernel individually. This is an enormously underconstrained attempt and leads to unpredictable behavior. CONSTSCALE assumes that β_f can be approximated by β_h , so $\tilde{\beta}_f$ must be an increasing function of β_g (see bottom panel of Fig. 5). According to Eq. (8), this ansatz is only applicable if β_g is negligible. For very small kernel scales, we can indeed see a tendency to converge to the correct solution, but for all other situations, this choice is manifestly non-optimal. Because of the clearly problematic behavior of CONSTSCALE and NMAX2, we exclude these two methods from the further investigation.

This situation is very similar for other choices of β_f and other PSF models. To work out the general trends of the three remaining methods, we average over all scales β_f and β_g and plot the results in dependence of the PSF model.

The top and middle panels of Fig. 6 confirm that all remaining methods yield essentially unbiased estimates of the shear, although we notice a mild tendency of SAME and SIGNIFIC to underestimate γ_1 . This indicates that truncation of the decomposition order $n_{\max}^h = n_{\max}^g$ might be insufficient for high S/N




Fig. 7. Analogous to Fig. 6 but for the high-noise simulations.

images. That this underestimation is absent at higher kernel orders confirms this interpretation.

Within the errors, the recovered scale size $\tilde{\beta}_f$ is rather unbiased (see the bottom panel of 6). For SAME and SIGNIFIC, we can see a clear shift of $\tilde{\beta}_f$ for PSFc. The reason for this lies in the large spatial extent and wide wings of the Airy disk model in combination with a low n_{\max}^h . Since the entries of P depend in a nonlinear way on β_f , this shift affects the recovery of the shear and leads to slightly poorer results.

From this initial simulation with moderate noise, we can conclude that one should respect Eq. (8) and must not truncate the shapelet series of h_n severely.

4.3. Performance with high noise

We now consider a realistic weak-lensing situation by increasing the noise level by a factor of 10, hence $4.5 \leq S/N \leq 22$ (cf. Fig. 3e). To partly compensate for the higher noise, we increase the number of realizations to $N = 1000$.

Considering Fig. 7, we can confirm that the shear estimates from these three methods are also not significantly biased for very noisy images. However, for SAME we can see a remarkable drop in the mean of γ_1 and a drastic increase in the noise in γ_1 and γ_2 with the kernel order. Both findings are probably related to using P^{-1} instead of P^\dagger when performing the deconvolution. In contrast to the two methods we are proposing here, SAME uses $\tilde{n}_{\max}^f = n_{\max}^g$ (cf. Table 1). For the typical weak-lensing scenario – characterized by $n_{\max}^f < n_{\max}^g$, where all methods create a substantial amount of overfitting, cf. Fig. 3f – this assumes finding a higher number of significant deconvolved coefficients than are actually available. These additional, noise-dominated coefficients affect Q_{ij} and γ_Q (cf. Eq. (16)). Therefore, these quantities




Fig. 8. Distance in shapelet space R_s between the mean deconvolved and the true intrinsic coefficients in dependence of the PSF model for the three methods FULL (top panel), SIGNIFIC (middle panel), and SAME (bottom panel). The mean is computed by averaging over all available values of β_f and β_g (in total 60 independent combinations). Shown are the results for the moderate-noise simulations (solid line) and the high-noise simulations (dotted line).

become rather noisy themselves. Given that those high-order coefficients contain mostly arbitrary pixel noise that does not have a preferred direction, they also tend to dilute the available shear information from the lower-order coefficients, which explains the drop in γ_1 . The estimate for γ_2 is not affected, as its true value was zero anyway.

The superior behavior of FULL and SIGNIFIC in these low S/N simulations can also be seen more directly. As measure of the decomposition quality, we calculate the distance in shapelet space between the mean deconvolved coefficients \tilde{f}_n and the true input coefficients f_n ,

$$R_s^2 = \sum_n (\langle \tilde{f}_n \rangle - f_n)^2. \quad (17)$$

Figure 8 confirms that, as long as the kernel order is small, all three methods perform quite similarly. But when the kernel order increases, SAME tries to recover a quadratically increasing number of deconvolved coefficients whose individual significance is lowered at the same time. On the other hand, FULL and SIGNIFIC make use of the redundancy of the overdetermined coefficient set, which is created by applying a rectangular matrix P in Eq. (6). As a direct consequence of computing the least-squares solution via P^\dagger , the higher the number of convolved coefficients and the lower the number of significant intrinsic coefficients, the better these intrinsic coefficients can be recovered from noisy




Fig. 9. Recovered shear $\tilde{\gamma}$ and intrinsic scale size $\tilde{\beta}_f$ (in units of the true scale size β_f) in the high-noise simulations in dependence on the S/N of the convolved galaxy. The binning (dotted lines) is defined by the octiles of the S/N distribution, therefore all bins contain the mean values of approx. 7 combinations of β_f, β_g for each PSF model, in total ≈ 35 independent settings. The data are plotted at the center of the bins and the methods are slightly offset horizontally for visualization purposes. Color code is as explained in Fig. 6.

measurements. This explains the decrease in R_s with the kernel order for these two methods.

However, FULL also does not perform perfectly. The bottom panel of Fig. 7 reveals a bias on $\tilde{\beta}_f$, independent of the PSF model. The reason for this is again overfitting. As FULL goes to higher orders than SIGNIFIC and SAME, it is even more affected by the pixel noise. As the decomposition determines β_h by minimizing χ^2 , β_h tends to become larger because this allows the model to fit a larger (increasingly noise-dominated) area, which reduces the overall residuals and thus χ^2 . SIGNIFIC and SAME behave similarly when the kernel order – and hence the decomposition order – becomes larger.

To prevent the shapelet models from creeping into the noisy areas around the object, it seems useful to constrain β_h not only from below but also from above. In addition to a guess on \tilde{n}_{\max}^f , we therefore impose a constraint $\beta_g < \beta_h < \sqrt{\beta_f^2 + \beta_g^2}$. Inferring both should be feasible when investigating observational data.

As our simulations comprise galaxy models of varying S/N^2 , it is illustrative to present the deconvolution results in S/N bins. Figure 9 confirms that the two methods we propose here are very robust against image degradation. This is remarkable, since many weak-lensing pipelines (and also SAME in this paper)

² The models for both f and g have unit flux, so the surface brightness of the convolved object h depends on β_f and β_g .

suffer from an underestimation of the shear, which becomes increasingly prominent with decreasing S/N (Massey et al. 2007). Our statement from above, in which we related this drop to the high number of insignificant coefficients obtained from a deconvolution using SAME, is further supported by this figure. It is obvious that – independent of the kernel model – a low S/N in pixel space results in a low S/N in shapelet space. By obtaining the least-squares solution for the f_n , FULL and SIGNIFIC boost the significance of the recovered coefficients and thus perform better in the low S/N regime. The reason why $\tilde{\gamma}_1$ from FULL is consistently but insignificantly lower than the estimates from SAME is still somewhat unclear. A possible reason is the generally higher number of shapelet coefficients h_n for FULL and thus a more noticeable noise contamination.

5. Conclusions

Based on an analytic consideration of shapelet convolution, we have studied algorithms for the PSF deconvolution of galaxy images in shapelet space. The starting point are sum rules for shapelet convolution, showing that the intrinsic shapelet orders of PSF and image add up in the convolved image and that the squares of their scales are also added. We suggest an algorithm respecting these sum rules as well as possible in the presence of noise, whose central step is the deconvolution of the convolved image with the pseudo-inverse of the convolution matrix. Applications to simulated images have shown that our algorithm performs very well and, in many cases noticeably better, than previously suggested methods. We identify three main reasons for the improved performance:

- As the sum rule for the shapelet order shows, the convolution transports power to higher shapelet modes. The mean signal-to-noise ratio of the convolved coefficients is thus reduced. Our reduction of the order during the deconvolution increases the signal-to-noise without the need for calibration.
- The sum rules typically require a much higher shapelet order than the χ^2 minimization, in particular if the PSF model is structured. Many of the high-order coefficients are thus highly insignificant. The significance of the coefficients is re-established by the reduction to the order \tilde{n}_{\max}^f , which is chosen such that the shapelet expansion contains only significant coefficients. Fortunately, \tilde{n}_{\max}^f mainly depends on the signal-to-noise ratio of the galaxy, and this dependence is rather weak. Many different galaxy shapes can be deconvolved with the same \tilde{n}_{\max}^f , so binning the galaxies into broad signal-to-noise bins will suffice.
- The STEP-2 project (Massey et al. 2007) has shown that shear measurements generally depend strongly on the galaxy brightness. Our algorithm seems to have the advantage of lowering the influence of pixel noise in a well-defined manner.

We shall proceed to study the performance of our algorithm in shear measurements under realistic conditions.

Acknowledgements. PM wants to thank Thomas Erben, Alex Boehmert, Ludovic van Waerbeke, and Marco Lombardi for very fruitful discussions. P.M. is supported by the DFG Priority Programme 1177. M.M. is supported by the Transregio-Sonderforschungsbereich TRR 33 of the DFG.

Appendix A: Convolution scale and order

For simplicity we restrict ourselves to the one-dimensional case. From Eqs. (1) and (5) it is apparent that

$$h(x) = \sum_{m,l} f_m g_l \int dx' B_m(x'; \beta_f) B_l(x - x'; \beta_g). \quad (\text{A.1})$$

We define $I_{m,l}(x; \beta_f, \beta_g)$ as the integral in Eq. (A.1) and decompose it into shapelets with scale size β_h and maximum order N ,

$$I_{m,l}(x; \beta_f, \beta_g) = \sum_n c_n B_n(x; \beta_h). \quad (\text{A.2})$$

Considering Eqs. (2), (3) and (A.1), we recognize that N cannot be infinite but is determined by the highest modes of the expansions of f and g , which we will call M and L , respectively. Restricting ourselves to these modes and dropping all unnecessary constants, we can proceed,

$$\begin{aligned} I_{M,L}(x; \beta_f, \beta_g) &= \\ & \int dx' (x')^M \exp\left[-\frac{x'^2}{2\beta_f^2}\right] (x - x')^L \exp\left[-\frac{(x - x')^2}{2\beta_g^2}\right] \\ &= \sum_{i=0}^L (-1)^{L+1} \binom{L}{i} x^{L-i} \int dx' (x')^{M+i} \exp\left[-\frac{(x - x')^2}{2\beta_g^2} - \frac{(x')^2}{2\beta_f^2}\right], \end{aligned} \quad (\text{A.3})$$

where we expanded $(x - x')^L$ in the last step. By employing Eq. (8) and substituting $\tilde{x} = x' - \frac{\beta_f^2}{\beta_h^2} x$, we can split the exponential,

$$\begin{aligned} I_{M,L}(x; \beta_f, \beta_g) &= \sum_{i=0}^L (-1)^{L+1} \binom{L}{i} x^{L-i} \exp\left[-\frac{x^2}{2\beta_h^2}\right] \\ & \times \int d\tilde{x} \left(\tilde{x} + \frac{\beta_f^2}{\beta_h^2} x\right)^{M+i} \exp\left[-\frac{\beta_h^2}{2\beta_f^2 \beta_g^2} \tilde{x}^2\right]. \end{aligned} \quad (\text{A.4})$$

Again, we expand $\left(\tilde{x} + \frac{\beta_f^2}{\beta_h^2} x\right)^{M+i}$, which yields the desired expression

$$\begin{aligned} I_{M,L}(x; \beta_f, \beta_g) &= \sum_{i=0}^L (-1)^{L+1} \binom{L}{i} \sum_{j=0}^{M+i} \frac{\beta_f^{2(M+i-j)}}{\beta_h^{M-L+2i-j}} \binom{M+i}{j} C_j \\ & \times \left(\frac{x}{\beta_h}\right)^{M+L-j} \exp\left[-\frac{x^2}{2\beta_h^2}\right], \end{aligned} \quad (\text{A.5})$$

where we inserted $C_j \equiv \int d\tilde{x} \tilde{x}^j \exp\left[-\frac{\beta_h^2}{2\beta_f^2 \beta_g^2} \tilde{x}^2\right]$. Apart from the omitted constants, the second line of A.5 is the definition of $B_{M+L-j}(x; \beta_h)$ (cf. Eqs. (2) and (3)), which shows that the natural choice is justified. Moreover, as j runs from 0 to $M+i$, we see that the maximum order N is indeed $M+L$, as we claimed in Eq. (9). Since $C_j = 0$ if j is odd, the only states with non-vanishing power have the same parity as $M+L$.

In summary, the natural choice is inherited from the Gaussian weighting function in Eq. (3) and the maximum order the result of a product of polynomials.

References

- Bartelmann, M., & Schneider, P. 2001, Phys. Rep., 340, 291
 Bergé, J. 2005, An introduction to shapelets based weak lensing image processing, 1st edn.
 Chang, T.-C., Refregier, A., & Helfand, D. J. 2004, ApJ, 617, 794
 Heymans, C., Van Waerbeke, L., Bacon, D., et al. 2006, MNRAS, 368, 1323
 Kuijken, K. 1999, A&A, 352, 355
 Kuijken, K. 2006, A&A, 456, 827
 Massey, R., Heymans, C., Bergé, J., et al. 2007, MNRAS, 376, 13
 Massey, R., & Refregier, A. 2005, MNRAS, 363, 197
 Melchior, P., Meneghetti, M., & Bartelmann, M. 2007, A&A, 463, 1215
 Refregier, A. 2003, MNRAS, 338, 35
 Refregier, A., & Bacon, D. 2003, MNRAS, 338, 48