Weak Lensing

David Wittman[1]

Bell Laboratories, Lucent Technologies, Room 1E-414,
700 Mountain Avenue, Murray Hill, NJ 07974, USA

Abstract. In the preceding chapters, the effects of lensing were so strong as to leave an unmistakable imprint on a specific source, allowing a detailed treatment. However, only the densest regions of the universe are able to provide such a spectacular lensing effect. To study more representative regions of the universe, we must examine large numbers of sources statistically. This is the domain of weak lensing.

1 Introduction

1.1 Motivation

Weak lensing enables the direct study of mass in the universe. Lensing, weak or strong, provides a more direct probe of mass than other methods which rely on astrophysical assumptions (*e.g.* hydrostatic equilibrium in a galaxy cluster) or proxies (*e.g.* the galaxy distribution), and can potentially access a more redshift-independent sample of structures than can methods which depend on emitted light with its r^{-2} falloff. But strong lensing can be applied only to the centers of very dense mass concentrations. Weak lensing, in contrast, can be applied to the vast majority of the universe. It provides a direct probe of most areas of already-known mass concentrations, and a way to discover and study new mass concentrations which could potentially be dark. With sources covering a broad redshift range, it also has the potential to probe structure along the line of sight.

Specifically, we might expect weak lensing to answer these questions:

- Where are the overdensities in the universe ?
- Are they associated with clusters and groups of galaxies ? Does light trace mass in these systems ?
- How much do these systems contribute to Ω_m, the mean density of matter in the universe ?
- What is their mass function and how does that function evolve with redshift ? What does that imply for the dark energy equation of state ?
- What are the structures on larger scales (walls, voids, filaments) ?
- Is this structure comparable to that seen in cosmological simulations? Which cosmology matches best ?
- What is the nature of dark matter ?
- Can observations of lensing put any constraints on alternative theories of gravity ?

Until recently, deep imaging on the scale required to answer the above questions with weak lensing was simply impractical. The development of large mosaics of CCDs has expanded the field greatly. The large data volume leads to ever-decreasing statistical errors, which means that very close attention must be paid to systematic errors and calibration issues. Weak lensing results must be carefully scrutinized and compared with those of other approaches with this in mind.

We start with a review of the basic concepts, the limits of weak lensing, and observational hurdles, and then address the above astrophysical questions.

1.2 Basics

The transition from strong to weak lensing can be seen at a glance in the simulation shown in Figure 1. Over most of the field, no one galaxy is obviously lensed, yet the galaxies have a slight tendency to be oriented tangentially to the lens. We seek to exploit this effect to derive information about the lens, and perhaps about the weakly lensed sources as well.

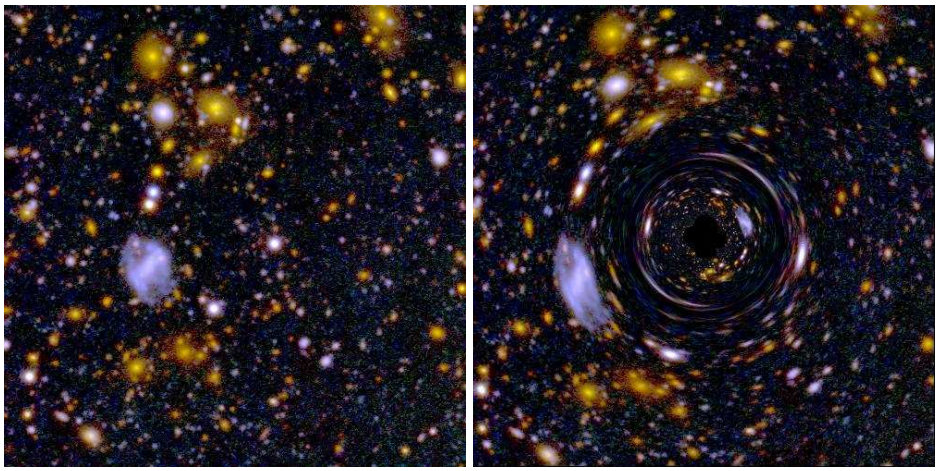

Fig. 1. Simulated effects of a lens: source plane (left) and image plane (right). Most regions of the lens can be probed only with weak lensing. Real sources are not in a plane, but this does not dramatically affect the appearance. Real lenses, such as galaxy clusters, would obscure much of the strong-lensing region.

We start with the *inverse magnification matrix* (see also Chapter on quasar lensing)

$$M^{-1} = (1-\kappa)\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \gamma \begin{pmatrix} \cos 2\phi & \sin 2\phi \\ \sin 2\phi & -\cos 2\phi \end{pmatrix}, \quad (1)$$

so called because it describes the change in source coordinates for an infinitesmal change in image coordinates, the inverse of the transformation undergone by the sources. This is Equation 16 of the Quasar Lensing chapter, which derives M^{-1} and defines the quantities within. We repeat here that the *convergence* κ represents an isotropic magnification, and the *shear* γ represents a stretching in the direction ϕ. They are both related to physical properties of the lens as linear combinations of derivatives of the deflection angle. However, κ can be interpreted very simply as the projected mass density Σ divided by the critical density $\Sigma_{\rm crit}$, while γ has no such straightforward interpretation. In fact, γ is *nonlocal*: its value at a given position on the sky depends on the mass distribution everywhere, not simply at that position. We will see this fact rear its ugly head in several places throughout this chapter. Shear is often written as a vector $\gamma_i = (\gamma\cos 2\phi, \gamma\sin 2\phi)$ or more succinctly as a complex quantity $\gamma e^{i2\phi}$.

Without multiple images of a source (as in the strong lensing case), we must have some independent knowledge of the sources if we are to measure magnification or shear. For example, if one source were a standard candle or ruler, the apparent magnitude or size of its image would immediately yield the magnification at that point. Of course, standard candles or rulers occur only in very special cases [17], so in practice we must analyze source *distributions*. We no longer get much information from a single source, and thus lose resolution; this is the tradeoff we must make for probing regions with weak tidal fields.

One source distribution that could be used in this way is $n(m)$, the number of galaxies as a function of apparent magnitude. In practice, this is difficult, because the measured slope of this distribution does not differ greatly from the critical slope at which equal numbers of galaxies are magnified into and out of a given magnitude bin, with no detectable change ($n \propto m^{0.4}$). There is enough difference to make some headway, but we would prefer to measure departures from zero rather than small changes in a large quantity.

The distribution of galaxy *shapes*, properly defined, does allow us to measure departures from zero. Approximate each source as an ellipse with position angle ϕ and (scalar) ellipticity $\epsilon = \frac{a^2-b^2}{a^2+b^2}$, where a and b are the semimajor and semiminor axes. Define a *vector ellipticity* $e_i = (\epsilon\cos 2\phi, \epsilon\sin 2\phi)$, or equivalently a *complex ellipticity* $\epsilon e^{i2\phi}$ (also called *polarization*). This encodes the position angle and scalar ellipticity into two quantities which are comparable to each other; the dependence on 2ϕ indicates invariance under rotation by 180°. Figure 2 gives a visual impression of ellipses in this space.

We can now quantify the visual impression of Figure 1. In the absence of lensing, as in the left panel, galaxies are randomly oriented: The observed distribution of e_i is roughly Gaussian with zero mean and an rms of $\sigma_e \sim 0.3$. In the presence of lensing, as in the right panel, this distribution is no longer centered on zero, as long as we consider an appropriately-sized patch of sky. In fact, we will *assume* that any departure from zero mean must be due to

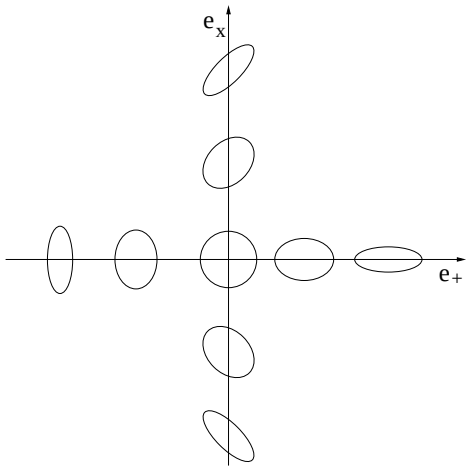

Fig. 2. A sequence of ellipses with various amounts of each ellipticity component. Inspired by the appearance of these ellipses, the two components are often labeled e_+ and $e_\times$.

lensing. We will examine the limits of this assumption in some detail later, but for now let us accept that on large enough scales, the cosmological principle demands it, and as a practical matter, we average over sources at a wide range of redshifts, which are too far apart physically to influence each other's alignment.

The effect of the magnification matrix on the complex ellipticity can be computed if M is constant over a source. This is obviously not valid for very large sources or those near caustics, but it is valid for the vast majority of the sky and for typical sources with sizes of a few arcseconds. The result is that $\epsilon^I = \epsilon^S + \frac{\gamma}{1-\kappa}$, where superscripts indicate image and source planes [18]. We don't know any of these quantities for a single source, but we do know (or assume for now) that $\langle \epsilon^S \rangle = 0$, where brackets indicate averaging over many sources. Hence

$$\langle \epsilon^I \rangle = \langle \frac{\gamma}{1-\kappa} \rangle. \tag{2}$$

The quantity on the right is called the *reduced shear* g. A second approximation we can often make is that $\kappa \ll 1$, so that $\langle \epsilon^I \rangle = \langle \gamma \rangle$. This is called the *weak lensing limit*.

The fundamental limit to the accuracy with which we can measure γ in the weak lensing limit is *shape noise*, or the width of the source ellipticity distribution $\sigma_e \sim 0.3$. Averaging over n sources should decrease the uncertainty to $\frac{\sigma_e}{\sqrt{n}}$, but n is limited by the depth of the observations and the area over which we are willing to average γ; these tradeoffs are discussed below. Also note that knowledge of the shear alone is not strictly enough to infer mass distributions because of the *mass sheet degeneracy* [44,111], introduced in a

different context in the Quasar Lensing chapter. This degeneracy arises because a uniform sheet of mass induces only magnification, not shear. Because the equations are linear, we could therefore add or subtract a mass sheet without affecting the shear. In practice, we can still answer many questions with shear alone, as discussed below.

1.3 Cosmology dependence

Both convergence and shear scale as the combination of angular diameter distances $\frac{D_{\rm LS}D_{\rm L}}{D_{\rm S}}$, or as the *distance ratio* $\frac{D_{\rm LS}}{D_{\rm S}}$ for a given lens. (Recall from the Quasar Lensing chapter that $D_{\rm LS}$, $D_{\rm L}$, and $D_{\rm S}$ are the angular diameter distances from lens to source, observer to lens, and observer to source, respectively. Note that $D_{\rm S} \neq D_{\rm L} + D_{\rm LS}$; see [59] for a quick review and [101] for a thorough treatment of distance measures in cosmology). This cosmology-dependent quantity is plotted as a function of source redshift in Figure 3 for several lens redshifts and two different cosmologies. In principle, this could be used to constrain the cosmology if source redshifts are known, and if the lens mass is known independently (the effects of a larger lens mass and a larger universe are degenerate). But this remains an unused cosmological test because lens parameters and source redshifts are usually poorly known. Usually, a cosmology is assumed and lens parameters are estimated using any available knowledge of source redshifts. Less often, a well-characterized lens is used to explore the source redshift distribution. However, source redshift distributions are usually quite broad, and weak lensing can only be used to estimate the *mean* distance ratio to a group of sources, which is not same as the distance ratio corresponding to the mean redshift. Section 1.4 deals with ways of estimating the mean distance ratio or otherwise accounting for a broad source redshift distribution.

Another way of viewing the same information is to fix the source redshift and plot this ratio as a function of lens redshift (Figure 4). This reveals the relative importance of different structures along the line of sight and is often called the *lensing kernel* or *lensing efficiency*.

1.4 Applicability of weak lensing

As with all astrophysical tools, we must be aware of the limitations of weak lensing before plunging into results. They include the weak lensing approximation itself; mass sheet degeneracy if only shear is used; poor angular resolution because of its statistical nature; source redshift difficulties; and possible departures from the assumption of randomly oriented sources. We now examine these limits and ways of dealing with them.

Weak lensing approximation The approximations that M is constant over each source and that $\kappa \ll 1$ cannot be applied when dealing with the

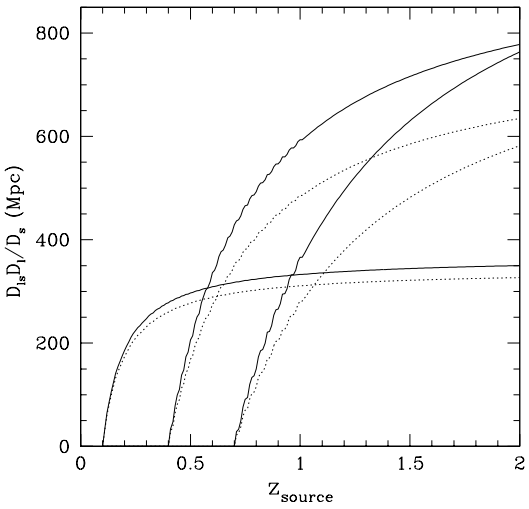

Fig. 3. $\frac{D_{\rm LS}D_{\rm L}}{D_{\rm S}}$ as a function of source redshift, for several lens redshifts (indicated by the intersections of the curves with the horizontal axis) and several cosmologies. The cosmologies are Λ-dominated (solid lines, $H_0 = 70$ km s^{-1} Mpc^{-1}, $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$) and open (dashed lines, $H_0 = 70$ km s^{-1} Mpc^{-1}, $\Omega_m = 0.4$, $\Omega_\Lambda = 0$). Each solid line is higher than its dashed counterpart, reflecting the larger size of the Λ-dominated universe. Although this quantity appears to be a sensitive test of the cosmology, it is degenerate with the lens mass.

centers of massive clusters and galaxies. Of course, analysis of such regions is not lacking—it is the topic of most of this book. Here we merely wish to mention work that has been done on combining weak and strong lensing information [1]. We also note that, where only the second approximation fails, Equation (2) can be solved iteratively for κ.

Mass sheet degeneracy Mass sheet degeneracy was a serious concern when fields of view were small and lens mass distributions extended well beyond the edges. Modern imagers now deliver fields of view $\sim 0.5^\circ$ on a side (> 3 Mpc radius for any lens at $z > 0.15$), so this concern has diminished. The degeneracy may also be broken by adding magnification information, which may come from strong lensing, or from a method called the depletion curve.

Magnification imposes two effects of opposite sign on the areal density of sources. Galaxies fainter than the detection limit (or any chosen brightness threshold) are amplified above the threshold, increasing the density of sources, but at the same time the angular separation between galaxies is stretched, decreasing the density of sources. The net effect depends on the slope of $n(m)$, the (unlensed) galaxy counts as a function of magnitude. A

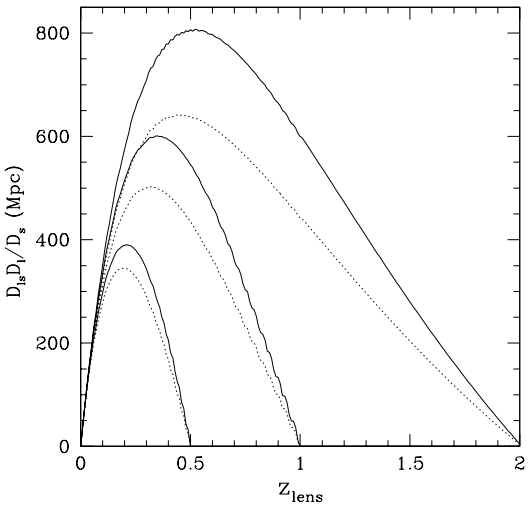

Fig. 4. Same as for Figure 3, but as a function of lens redshift, for several values of source redshift (which correspond to the right-hand end of each curve). The lensing efficiency is a very broad function, making it difficult to separate unrelated structures along the line of sight.

logarithmic slope less than 0.4 (usually the case at visible wavelengths, but barely) will not provide enough "new" sources to overcome the dilution effect, so the source density decreases as κ increases toward the center of a cluster. This *depletion curve* reveals lens parameters, as shown in Figure 5 [88]. Despite the name, the method need not be restricted to one-dimensional information [22]; [88] includes a lens ellipticity and position angle estimate based on a crude depletion map. In practice, measuring magnification is quite difficult, because the slope of $n(m)$ is perilously close to 0.4, and there are few published depletion curve measurements [88,40]. For the remainder of this work, we shall concentrate on algorithms and results using shear, not magnification.

Angular resolution The angular resolution of weak lensing is limited by the areal density of sources. With a shape noise of $\sigma_e \sim 0.3$ and $\sqrt{n}$ statistics, a shear measurement accurate to p percent requires $\sim 1000p^{-2}$ sources. Angular resolution is then set by the area of sky over which these sources are scattered. This in turn depends on the depth and wavelength of the observations; in R there is one source per square arcminute in a one-magnitude wide bin at $R \sim 21.4$, increasing by a factor of ~ 2.5 for every magnitude deeper [122]. A medium deep observation capable of shape measurements to $R \sim 25$ thus yields about 20 galaxies arcmin^{-2} (assuming a bright cutoff $R > 23.5$ to

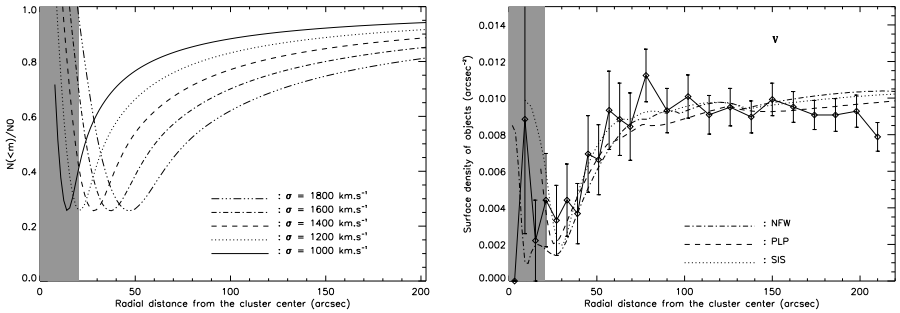

Fig. 5. Left: theoretical depletion curves for a variety of lens velocity dispersions (lens mass $\propto \sigma_v^2$). Right: depletion curve observed for MS1008-1224 in V band. From [88].

eliminate largely foreground sources), implying that 2 arcmin^2 are required for 5% accuracy in shear.

Getting more sources per unit area requires much more telescope time. Source density will ultimately be limited by confusion — when sources are so numerous that they overlap and hinder shape measurements — around $\sim$ 1000 sources arcmin^{-2} for ground-based data. This implies $\sim 20''$ shear resolution, or better for space-based data, if galaxy counts keep rising at the same rate. However, such depth is hard to come by and must compete against area and wavelength coverage (useful for constraining source redshifts) when planning for a given amount of telescope time.

Another tradeoff commonly used is to sacrifice resolution in one dimension to achieve better resolution in the other. Clusters are commonly analyzed in terms of a radial profile, which assumes they are axisymmetric and allows all sources at a given radius from the cluster center to be averaged together. Less massive clusters and groups can be "stacked" to yield an average profile with reasonable resolution, just as in galaxy-galaxy lensing [57,116].

Source redshift distribution Lack of knowledge of the source redshift distribution is often a limit in calibrating weak lensing measurements. The root of this problem is that deep imaging quickly outruns the ability of even the largest telescopes to provide spectroscopic redshifts for a fair sample of sources.

The recent development of photometric redshift techniques, in which multicolor imaging provides enough spectral information for a reasonable redshift estimate citeConnolly1995,Hogg1998, has brought hope that source redshifts may be estimated to sufficient accuracy from imaging alone. For example, the Hubble Deep Field yielded photometric redshifts accurate to ~ 0.1 per galaxy in the redshift range $0 - 1.4$ with seven filters extending through the near-infrared ($UBVIJHK$) [60]. A look at Figure 3 shows that this provides

a reasonable accuracy in distance ratio in most situations. The accuracy improves with the number of filters used, resulting in a tradeoff between accuracy and telescope time. Deep U and infrared imaging are much more expensive than $BVRI$ in terms of telescope time, but it is difficult to effectively cover a large redshift range with only $BVRI$. Few spectral features are to be found in the observed $BVRI$ bandpasses for sources in the redshift range $\sim 1.5 - 3$, which greatly increases uncertainties there.

However, these problems are not fundamental, and photometric redshifts will become routine. They will do much more than help estimate the mean distance ratio required for calibrating lenses. Because sources lie at a range of redshifts, they will provide the opportunity to probe structure along the line of sight (albeit with resolution limited by the width of the lensing kernel). The ultimate goal is *tomography* — building up a three-dimensional view of mass in the universe from a series of two-dimensional views at different redshifts. The combination of weak lensing and photometric redshifts thus promises to be very powerful, but as yet there are not many published examples of combining the two, and little theoretical work on optimal ways of doing so. Although we can expect photometric redshifts to be a routine part of future lensing work, we must be aware of alternative ways of confronting the source redshift problem.

First, some questions can be answered without calibration of source redshifts. The two-dimensional morphology of a cluster lens is one example — the source redshift distribution should not depend on position (as long as magnification is negligible and cluster members do not contaminate the source sample). Similarly, source redshifts are not required for discovery of mass concentrations in surveys, but without them, the volume probed is unknown. Clearly, the questions which can be answered this way are limited.

A more general calibration strategy is through additional, identical observations of a "control lens" of known redshift and mass (*e.g.* a cluster with a dynamical, X-ray, and/or strong lensing mass estimate). This does allow estimation of the mean distance ratio to a population of sources much too faint to reach with spectroscopy, but it certainly has its limits. It is difficult to obtain identical observations, and the (probably considerable) uncertainty in the mass of the control lens becomes a systematic for the rest of the data. But more fundamentally, shear from the control lens samples only that part of the source distribution which is behind the control lens, so that strictly speaking, a control lens must be at the same redshift as the target. For weak lensing by large-scale structure, the distribution, not simply the mean distance ratio, is required. This would require control lenses at a range of redshifts, which is impractical. Photometric redshifts should do a much better job with more realistic data requirements. Even in the age of photometric redshifts, though, this method will have its role. The shear induced by calibrated lenses will provide a check on photometric redshift estimates, which may not be checkable with spectroscopy if applied to very faint sources.

Another strategy is keeping the imaging as shallow as current redshift surveys, which go to $R \sim 24$. One can then look up the median redshift for any magnitude cut; for $23 < R < 24$, for example, the median redshift is 0.8 [29]. Even the redshift distribution is known to some extent, with 120 sources in that magnitude slice in the survey cited. Shallow imaging need not probe a small volume, as a large area can be covered with a reasonable amount of telescope time. But it does limit the distance probed and the angular resolution of the mass reconstruction (because the areal density of sources is low at $R \leq 24$). This strategy also limits selection of sources based on color, which is very useful for limiting contamination by galaxies in a cluster being studied (or for de-emphasizing foreground contamination in general) because the median redshift of a color-selected sample is not yet something that can be looked up in a redshift survey.

Intrinsic alignments The crucial assumption in weak lensing is that the sources have random intrinsic orientations, so that any departure from randomness is due to lensing. This assumption is worth examining before proceeding further. We will concentrate on potential damage to measurements of weak lensing by large-scale structure (cosmic shear), because the lensing signal from clusters is usually at a much higher level. However, it is worth keeping in mind that all applications of weak lensing could be affected at some level.

The first detections of cosmic shear in 2000 motivated several analytic [25] and computational [54,36] studies of intrinsic alignment mechanisms, and the field is still sorting itself out. There are several mechanisms which could produce such intrinsic alignments, including tidal stretching of galaxies in a gravitational potential, and coupling of the potential to the spin vectors of galaxies [34]. The amount of alignment predicted as a function of angular scale varies greatly depending on the mechanism and the strength of the coupling; it remains unknown which model, if any, is correct. However, in most scenarios, intrinsic alignments would represent a $\lesssim 10\%$ contamination of the cosmic shear measurements.

While the situation is still evolving, one rule is certain: the effect of any intrinsic alignment is diluted when sources lie at a large range of redshifts, as is naturally the case in deep imaging. As we shall see, the signal from lensing by large-scale structure increases with source redshift. Hence, lensing must dominate at high enough source redshift, and intrinsic alignments at low enough source redshift. This is illustrated by Figure 6, which shows predicted intrinsic alignment and cosmic shear levels for several source redshifts. For shallow surveys like the Sloan Digital Sky Survey [141] (SDSS), intrinsic alignment may frustrate attempts to measure cosmic shear, but deeper surveys specifically designed to measure cosmic shear are safe. Indeed, the roughly one million spectroscopic redshifts SDSS plans to acquire will be invaluable in measuring intrinsic alignments precisely, and their measurements,

after scaling to higher redshifts, in turn may facilitate estimation and even removal of intrinsic alignment effects from the deeper surveys. Deep surveys may also be able to provide a lensing signal using only sources which cannot be physically associated, as indicated by their photometric redshifts. Density reconstruction methods in the presence of intrinsic alignments are already being investigated [83].

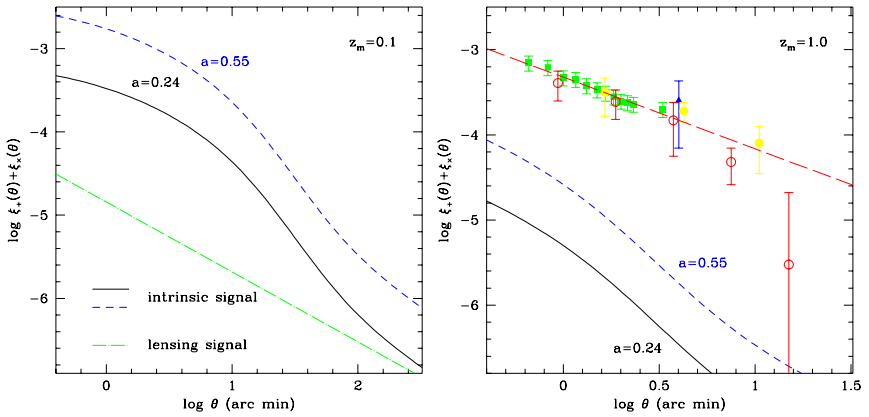

Fig. 6. The importance of intrinsic alignments depends strongly on source redshift. The expected levels of ellipticity correlation (defined in Section 3.1) due to intrinsic alignments and to weak lensing by large-scale structure are shown for low-redshift (median source redshift $z_m = 0.1$) surveys in the left panel and for high-redshift ($z_m = 1$) surveys in the right panel. In each case, the straight line indicates the expected signal from weak lensing, and the curves indicate the expected signal from intrinsic alignments, for two different values of a spin-coupling parameter, giving some idea of the modeling uncertainties. The right panel also contains cosmic shear measurements from the literature, which all happen to have $z_m \sim 1$. Adapted from [34].

The two or three detections of intrinsic alignments in real data are indeed at low redshift. Ellipticity correlations have been reported in SuperCOSMOS data [23], but because there is no redshift information, intrinsic alignments can only be inferred (the median source redshift is estimated to be < 0.1, so that the inferred cause is intrinsic alignments rather than lensing). Spin alignments have been found in the Tully catalog, which consists of several thousand nearby (within a few Mpc) spirals [102], and in the PSCz, a redshift survey of 15500 galaxies detected by the IRAS infrared satellite mission [84]. Because these catalogs have redshift information, these represent solid detections. However, spin correlations are one step removed from ellipticity correlations, which are the relevant quantity for lensing.

It may also be possible to extract intrinsic alignments from the lensing data itself. To first order, lensing produces a curl-free, or E-type (in analogy with electromagnetism) shear field. (Multiple scattering can produce a weak divergence-free, or B-type field, but that can be safely ignored for the moment.) Therefore, decomposition of a measured shear field into E-type and B-type fields might allow separation of the lensing and intrinsic alignment effects [35]. This decomposition is difficult, but a crude indicator of the B-type field is the traditional 45° test. In this test, a lensing signal should vanish when one component of the shear is exchanged for the other (equivalent to rotating each source by 45°), and nonzero results would indicate a systematic error. All published cosmic shear results were vetted using this test among others, with no indication of contamination. However, not all intrinsic alignment mechanisms produce B-type power; an example is tidal stretching (tidal fields are the basic mechanism for both stretching and lensing, after all). Based on other astrophysical arguments, tidal stretching is not likely to be significant [35], but even angular momentum coupling models can produce much more E-type than B-type correlations [86]. Passing the 45° test is a necessary but not sufficient condition for confidence in the results.

In summary, intrinsic alignments are not to be dismissed. They must be addressed and may even dominate the lensing signal in certain low-redshift applications. However, the dilution effect of a broad source redshift distribution means that none of the conclusions of weak lensing to this point can be called into doubt. Ongoing and future weak lensing studies may have to apply small corrections for this effect, but how small is still uncertain. Accurate corrections will probably be available by the completion of the SDSS, which will do much to increase our knowledge of intrinsic alignments in the nearby universe.

1.5 Measuring shear

In most weak lensing work, a source galaxy is approximated as an ellipse fully described by its quadrupole moments

$$\begin{aligned} I_{xx} &\equiv \frac{\Sigma I w x^2}{\Sigma I w} \\ I_{yy} &\equiv \frac{\Sigma I w y^2}{\Sigma I w} \\ I_{xy} &\equiv \frac{\Sigma I w x y}{\Sigma I w} \end{aligned} \tag{3}$$

where $I(x, y)$ is the intensity distribution above the night sky level, $w(x, y)$ is a weight function, the sum is over a contiguous set of pixels defined as belonging to the galaxy, and the coordinate system has been translated so that the first moments vanish (*i.e.* the centroid of the galaxy is chosen to be the origin of the coordinate system). Early work used intensity-weighted moments ($w = 1$),

but it was realized that this produces ellipticity measurements with noise properties that are far from optimal or even divergent. Now, w is usually chosen to be a circular [73] or elliptical [15] Gaussian, which deweights the outer pixels which have a big lever arm but low signal-to-noise. The two ellipticity components can be defined as [130]

$$\begin{aligned} e_+ &= \frac{I_{xx} - I_{yy}}{I_{xx} + I_{yy}} \\ e_\times &= \frac{2I_{xy}}{I_{xx} + I_{yy}} \end{aligned} \tag{4}$$

These are related to the scalar ellipticity ϵ and position angle ϕ by

$$\begin{aligned} \epsilon &= (e_+^2 + e_\times^2)^{\frac{1}{2}} \\ \phi &= \frac{1}{2} tan^{-1}(\frac{e_\times}{e_+}) \end{aligned} \tag{5}$$

Then a simple estimate of the shear in the weak lensing limit is $\gamma_i = \langle e_i \rangle / 2$, where the brackets denote averaging over many sources (perhaps with weighting of the sources based on estimated measurement errors, redshift, etc.) to beat down shape noise. Note that this definition of ellipticity differs from that in Equation (2) by a factor of two; both definitions are presented here because both are common in the literature. This latter estimator sometimes called the *distortion* statistic. Also, there are alternative formulations in terms of octupole moments [48], Laguerre expansions [15] and shapelets [107,108,27]. Before applying any of these estimators, we must account for the effects of point-spread function (PSF) anisotropy and broadening.

PSF anisotropy No optical system is perfect, and PSFs on real telescopes tend to be $\sim 1 - 10\%$ elliptical. This constitutes a huge systematic error, of the order of the shear induced by even a massive cluster, and it must be removed as completely as possible before analyzing any galaxy shapes, and monitored afterward. The removal can be done *after* measuring shapes, by essentially subtracting the moments of the PSF from the galaxy moments, but a more computationally stable method is to remove this effect from the image, *before* measuring shapes. This is done by convolving the image with a kernel with ellipticity components opposite to that of the PSF [47]. The raw PSF is almost certainly position-dependent; therefore the circularizing kernel is also, but the convolved PSF is everywhere round. A round PSF is called *isotropic*, but keep in mind that this does not imply homogeneous: The convolved PSF may vary somewhat in size, because of the position dependence of the original PSF and of the small broadening introduced by the kernel. A more sophisticated scheme would introduce more broadening in the right places, leading to a PSF which is homogeneous as well as isotropic. It is also possible to choose a sharpening kernel, but this would amplify the noise in the image.

Figures 7 and 8 illustrate the effectiveness of the convolution procedure. Although there are low-level residuals in the convolved image, their lack of spatial correlation means that they will have difficulty masquerading as a weak lensing effect. Note that these PSF anisotropies change with time, as telescope temperature, focus, and guiding drift, so that each exposure must be treated separately. A possible benefit here is that if the anisotropies are really uncorrelated temporally, coaddition of multiple exposures will beat down the shape errors. Also, each CCD in a mosaic must be treated separately, as some discontinuities may arise from small differences in piston between devices.

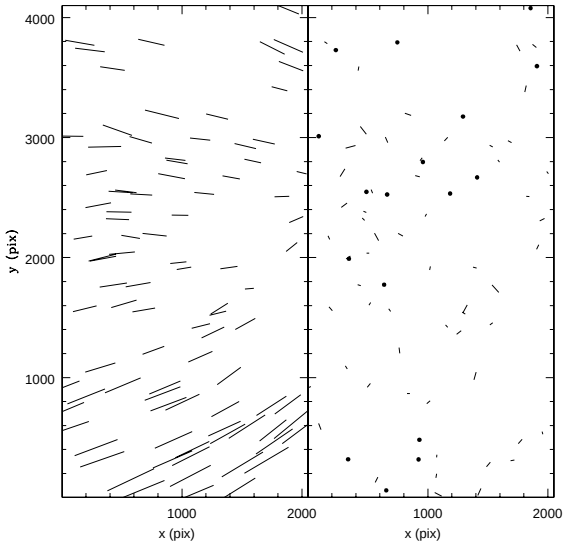

Fig. 7. Point-spread function correction in one 2k×4k CCD. Shapes of stars, which as point sources should be perfectly round, are represented as sticks encoding ellipticity and position angle. Left panel: raw data with spatially varying PSF ellipticities up to 10%. Right panel: after convolution with a spatially varying asymmetric kernel, ellipticities are vastly reduced (stars with $\epsilon < 0.5\%$ are shown as dots), and the residuals are not spatially correlated as a lensing signal would be.

PSF broadening Any effect which broadens the PSF will reduce the measured ellipticities of source galaxies which are not much larger than the PSF—generally including the distant galaxies most appropriate for lensing—because they will be broadened relatively more along their minor axes than along their major axes. In ground-based data, the dominant effect is "seeing", the broadening of the point-spread function due to turbulence in the atmosphere. Note that this is distinct from PSF anisotropy, which is caused by the telescope

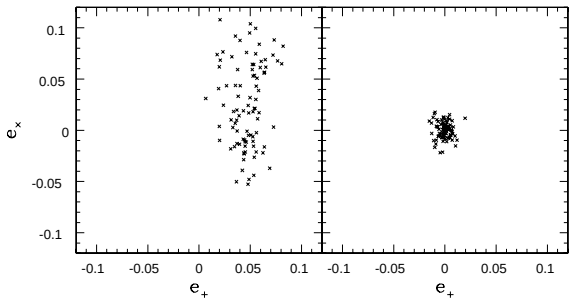

Fig. 8. Another way of plotting the efficacy of the PSF correction, often seen in the literature. For the same dataset shown in Figure 7, the ellipticity components of point sources are shown in a scatterplot, before and after correction. This type of plot hides any spatial correlation which may exist among the residuals, but Figure 7 shows that the residuals are uncorrelated in this case.

and camera optics. In fact, seeing produces a circular PSF as long as the integration time is much longer than the coherence time of the atmosphere (very roughly 30 ms at visible wavelengths); anyone who has observed in terrible seeing has probably noticed that at least the PSF is nicely round ! The effects of PSF anisotropy and broadening are sometimes called "shearing" and "smearing", respectively. The former has the effect of introducing a spurious weak lensing signal if uncorrected, while the latter has the effect of reducing any weak lensing signal.

There are several ways of correcting for smearing. The first is measuring the dilution of a simulated weak lensing signal relative to an unsmeared image, either simulated or perhaps from the Hubble Deep Fields. The seeing-free images are sheared by a known amount, convolved with the point-spread function of the real data, repixelized, and the shear measured. This has the advantage of including some effects which cannot be accounted for analytically, such as the coalescing of separate objects into an apparently elliptical single object.

On the other hand, a global correction is rejected by those who prefer to tailor the corrections to individual sources; after all, a large galaxy is smeared relatively less by seeing than is a small one. The advocates of this approach tend to use the KSB method, an analytical approach which takes into account the size of the PSF and of each source [73]. The KSB method also accounts for PSF shearing, but it can just as well be applied to a convolved image. See [70,71,82,15] for limitations of and possible successors to the KSB method. These approaches also weight each source according to its ellipticity uncertainty when computing a shear estimate.

Source selection Not every source in a deep image should be included in a shear measurement. A typical deep image includes stars, other unresolved sources, foreground galaxies, cluster members if the target lens is a cluster, and spurious objects, such as bits of scattered light around very bright stars. Getting rid of these unwanted sources is something of an art, which must reflect the particular data set, but generally there are four kinds of cuts. Magnitude cuts help get rid of stars (for reasonable galactic latitude, stars outnumber galaxies for $R \lesssim 22$ while galaxies greatly outnumber stars for $R > 23$) and bright foreground galaxies. Galaxies have a broad luminosity function, so such a cut is never completely effective at eliminating the foreground, but it helps. Color cuts seek to emphasize the faint blue galaxies at $z \sim 1$ [124]. If the target is a cluster, the cut should be blueward of the cluster's color-magnitude ridge. Even so, some cluster members and other foreground galaxies will survive. Size cuts eliminate unresolved objects, which at the relevant magnitudes include some stars, but mostly unresolved galaxies. Finally, cuts designed to insure that an object is not spurious must depend on the type of data available. Examples include rejecting objects that appear on only one of a multicolor set of images, and rejecting high ellipticity objects which are likely to be unsplit superpositions of two different objects.

Sanity checks There are a number of sanity checks that should be performed before believing any weak lensing result. In addition to the 45° test mentioned above, randomizing source positions while retaining their shapes should result in zero signal. Another good sanity check is correlating the source shapes with an unlensed control population, such as a set of stars. Finally, there are checks on the basic integrity of the catalog, such as the position angle distribution of sources, which might reveal spurious objects aligned with the detector axes. Because setting the source selection criteria can be somewhat subjective, it is also good to check that the results do not depend crucially on the exact magnitude or color cut.

2 Lensing by clusters and groups

Clusters of galaxies have long been studied from two somewhat opposing points of view. Visible from great distances, they are a convenient tracer of structure in the universe back to roughly half its present age. When examined individually, they are interesting astrophysical laboratories in their own right, with a variety of physical conditions and histories. But if so, they cannot be simple cosmological probes. So the study of clusters as astrophysical laboratories must inform and refine the study of clusters as cosmological probes.

What lensing adds to the study of clusters is a direct mass measurement without any assumptions about the dynamical state of the cluster. The first clusters were "weighed" in the 1930's with the dynamical method—assuming

that clusters are in virial equilibrium, the virialized mass is easily computed from the velocity dispersion. In the late 1980's, X-ray imaging of hot intracluster gas began to provide mass estimates, assuming hydrostatic equilibrium. In the 1990's, lensing began to provide mass estimates free of any such assumptions. The frequent agreement of the three types of estimate indicates that the dynamical assumptions are often valid, but the exceptions need to be identified. Those exceptions must be discarded from any samples used as cosmological probes, but they are often studied more closely for what they might reveal about mergers or other nonequilibrium processes.

In the past decade, it was enough simply to compare lensing measurements of cluster masses with those provided by other techniques. Driven by advances in wide-field detectors, we can now use lensing to *search* for clusters (or at least mass concentrations), and even estimate their redshifts. Shear-selected samples of clusters, free of any bias toward baryons that optically and X-ray selected samples might have, are currently being compiled. Comparison of the different types of samples will be instructive, either by confirming the use of traditional baryon-selected samples as cosmic probes, or perhaps by providing some counterexamples.

2.1 Masses and profiles

The first evidence of lensing by clusters came in the late 1980's in the form of strongly lensed giant arcs [85,118], which were used to constrain the mass inside the radius at which the arcs appeared. This was soon extended to somewhat less strongly lensed "arclets", and by 1990, to the first detection of what we now call weak lensing, the coherent alignment of thousands of weakly lensed background galaxies [125]. This alignment was measured in terms of the *tangential shear* γ_t, which is the component of shear directed tangential to an imaginary circle centered on the cluster and running through the source. The tangential ellipticity of a source is $e_t \equiv \epsilon\cos(2\theta)$ where θ is the angle from the tangent to the major axis of the source (Figure 9). When computing e_t from the ellipticity components, use the rotation

$$\begin{aligned} e_t &= +e_+\cos(2\beta) + e_\times\sin(2\beta) \\ e_c &= -e_+\sin(2\beta) + e_\times\cos(2\beta) \end{aligned} \tag{6}$$

where the angle β is also shown in Figure 9. Here e_c is a control statistic measuring the alignment along an axis 45 degrees from the tangent, which is not affected by an axisymmetric lens.

Methods for constraining cluster masses using tangential shear followed soon after the first detection of the effect [91]. The most important of these is *aperture densitometry*, which relates γ_t to the difference between the mean projected mass density inside a radius r_1 and that between r_1 and a larger radius r_2 [45]:

$$\bar{\kappa}(<r) - \bar{\kappa}(r_1 < r < r_2) = \frac{2}{1-r_1^2/r_2^2}\int_{r_1}^{r_2}\frac{\gamma_t}{1-\kappa(r)}d\ln r. \tag{7}$$

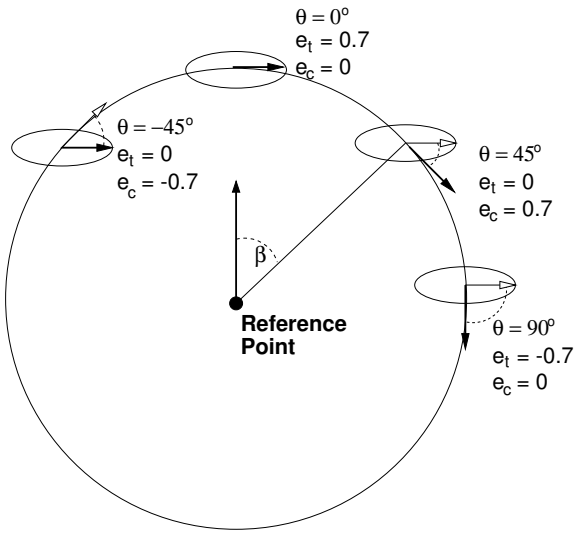

Fig. 9. Tangential ellipticity e_t of an $\epsilon = 0.7$ source with respect to a reference point. e_t carries the lensing signal, and the 45-degree component e_c serves as a control. In the presence of lensing but not shape noise, all sources would have $e_t > 0$ and $e_c = 0$; in practice $\langle e_t \rangle > 0$ and $\langle e_c \rangle = 0$.

The factor $1 - \kappa$ can be ignored in the weak lensing limit, but for massive clusters may not be ignored, leading to an iterative solution for κ (*e.g.* [47]). The left hand side of this equation is sometimes called the *zeta statistic* $\zeta(r_1, r_2)$. Note that this formula makes the mass sheet degeneracy explicit by specifying only relative values of κ; the best that can be done is extend r_2 to a very large value, at which κ should vanish.

A profile can be built up by repeatedly applying this statistic at a sequence of different r_1. However, this makes the points in the profile dependent on each other, as they use much the same data. If the goal is to find the best fit of a given type of profile, it is simpler to compute γ_t in a series of independent annuli and fit the shear profile expected from the mass model straightforwardly with least-squares fitting.

Weak lensing mass profiles are usually well fit by a singular isothermal sphere (SIS) or Navarro-Frenk-White (NFW, [98,99]) profile (see [76] for an extensive list of profiles used in lensing, along with their associated formulae). However, the nature of weak lensing makes it difficult to distinguish between models on two accounts. First, shear profiles do not have good dynamic range because the uncertainty in shear measurements increases dramatically at small radii, where there are not enough sources to beat down the shape noise. This is illustrated in the top panel of Figure 10. Note that this figure is for a very massive cluster; the signal-to-noise ratio can only be lower for less-massive clusters. Second, mass profiles which differ significantly

inside the radius where shear is measured can produce shear profiles which differ significantly *only outside* that radius. This is illustrated in the middle panel of Figure 10. The ability to distinguish between NFW and SIS (or more generally, power-law profiles) thus depends strongly on the size of the field [77], but Figure 10 demonstrates that a significant ambiguity remains even with a state-of-the-art imager with a 35' field. Weak lensing is therefore not definitively revealing cluster mass profiles, as one might have expected. Progress toward larger fields will be slow, as most large telescopes already have imagers which fill their usable fields of view. More likely, progress will come by adding magnification information. Finally, note that the most active (and revealing of the nature of dark matter) controversy surrounding cluster profiles involves cuspiness at the center, and this is not well addressed by weak lensing, with its poor angular resolution.

Initially, weak lensing analyses of clusters concentrated exclusively on the most massive clusters which were guaranteed to give a good signal. As the technique has matured, it has been extended to less massive, but more typical clusters [139]. The ultimate extension has been to groups; although a group by itself does not provide enough shear to get an accurate mass estimate, they can be "stacked" as in galaxy-galaxy lensing to build up an average profile with reasonable signal-to-noise [57]. The idea of stacking to obtain a good estimate of the average profile has been used for typical clusters as well [116]. Caution is required when interpreting "average" results, though, because they may be biased by a few unrepresentative systems, or in the worst case, meaningless if the sample is sufficiently heterogeneous.

Mass estimates of clusters and groups derived from weak lensing generally agree with estimates from velocity dispersions and X-ray imaging (see [89] for a list of published mass estimates as of 1999; there are now too many to list). At one time there was an apparent systematic discrepancy between (strong) lensing and X-ray estimates [93], but it was shown to be due to the complexities of the X-ray-emitting gas dynamics [3]. Hydrostatic equilibrium alone was shown to be less constraining than initially thought; temperature maps were needed [8]. XMM and Chandra now provide these, along with vastly improved angular resolution which allows for better treatment of cooling flows, and the first results show good agreement with lensing [4]. Of course, not every cluster behaves so well, and when there is disagreement a closer look often reveals interesting astrophysics such as cooling flows, mergers and their associated "cold fronts" or shock heating. In fact, $\sim$50% of clusters show some substructure in the X-ray [114]. Weak lensing can still supply the total mass, but due to its poor angular and line-of-sight resolution, the detailed work of disentangling the structures must be left to X-ray and dynamical measurements.

The three approaches, after all, have some fundamental differences which are not often mentioned. Dynamical estimates based on the virial theorem measure the total virialized mass of the cluster, while lensing can only mea-

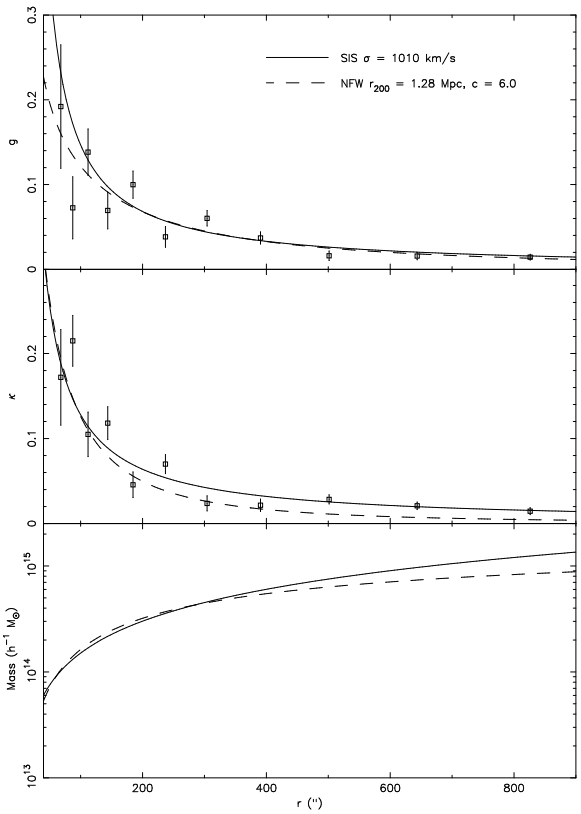

Fig. 10. Comparison of shear, convergence, and mass profiles for the massive cluster Abell 1689. *Top:* (Reduced) shear profile, with best-fit SIS and NFW profiles almost indistinguishable. *Middle:* κ profiles, showing that the NFW model falls off much more steeply than the SIS at large radii. This could barely be seen in the shear profile because of the nonlocality of shear. Despite appearances, the data do not favor the SIS. The points here are plotted assuming that outside the largest radius measured, the shear falls off as an SIS. If NFW is instead assumed for radii outside the measurement area, the points at large radius fall significantly. *Bottom:* Enclosed mass profiles of the two models. The physical scale is roughly 2 kpc arcsec^{-1}. From [28].

sure mass projected inside a certain radius. Even with a simplifying assumption such as spherical symmetry, a fair comparison of the two is difficult. Virial masses go as the square of the velocity dispersion, so small-number statistics and outliers can have a large effect on the mass [109]. Mass estimates from a dozen members may be good for a back-of-the envelope comparison, but beyond that should be treated with extreme caution. Velocity *profiles* would be more comparable to lensing and X-ray data. These are available for few clusters, but thanks to multi-object spectrographs on large telescopes, such

detailed dynamical analyses are becoming more common [16]. X-ray emission is proportional to the square of the density, so it is more sensitive to substructure than is lensing, which is simply proportional to the density. A fourth approach, the Sunyaev-Zel'dovich effect (SZE), measures the decrement in the cosmic microwave background (CMB) caused by upscattering of CMB photons by the hot intracluster gas. Like lensing, it is proportional to the density, but like X-ray emission, it depends on the density of baryons, rather than all matter. The first SZE measurements are starting to arrive and will soon offer their unique point of view.

Lensing stands out from X-ray and dynamical methods in being a projected statistic, so it is worth asking whether this introduces any bias. It appears that anisotropy in simulated clusters has little systematic effect [21], and so do uncorrelated structures along the line of sight [56]; both effects are around the 5% level. However, in reality there are also correlated structures along the line of sight, and these can bias masses upwards by tens of percent [26,90]. There is general agreement that the effect of other structures along the line of sight increases with aperture size. The bias changes with redshift in two ways. First, it is minimized when the cluster is near the peak of the lensing kernel (Figure 4), because other structures will be deemphasized. Second, younger clusters may have more nearby material, although this effect has not been investigated thoroughly [90]. Finally, note that the mass function is susceptible to bias even when an estimator is unbiased but has scatter, because there are more low-mass clusters to be scattered up than high-mass clusters to be scattered down (this applies equally to other types of mass estimates such as dynamical and X-ray) [90].

2.2 Two-dimensional structure

From the first detection of weak lensing, it was realized that the tangential shear procedure could be repeated about any reference point, not only the cluster center. By repeating it at a grid of points, a two-dimensional "mass map" (actually a map of κ) was constructed [125]. This was soon put on a firm theoretical footing by the derivation of a relationship between the Fourier transforms of κ and γ, starting from their relationship as different linear combinations of the second derivatives of the lensing potential ψ [72] (see the Quasar Lensing chapter for the relationship between ψ and κ). Essentially, κ can be expressed as a convolution of γ over the entire plane (there is also a real-space equivalent [47]). Of course, observations do not cover the entire plane, a problem called the *finite field effect* [10]. This is another manifestation of mass sheet degeneracy, as the spatial variation, but not the mean value, of κ can be reconstructed. Several reconstruction methods based on magnifications have been proposed to combat this problem [22,11], but, as mentioned above, magnification is very difficult to measure, and these methods have not been widely used. To a large extent, technology has solved the

problem, at least for clusters, by providing ever-wider fields of view, at the edges of which κ is presumably negligible.

In addition to direct *reconstruction* methods, there are *inversion* methods which solve for ψ, from which κ can be derived [12]. An extensive comparison of different methods found none to be clearly superior [119], although inversion methods tend to make it easier to include additional constraints such as those from magnification measurements or strong lensing features.

Many clusters have now been mapped using these techniques, and the mass distributions recovered are generally not surprising. That is, they roughly follow the optical and X-ray light distributions, on the scales which weak lensing is able to resolve. A vivid example of two-dimensional mass reconstruction is that of the supercluster MS0302+17 [74] (Figure 11). This supercluster contains three clusters separated by 15-20′ on the sky ($\sim 3-4$ Mpc transverse separation at a redshift of 0.42). All three clusters are recovered by the reconstruction algorithm, above the level of other, presumably noise, peaks.

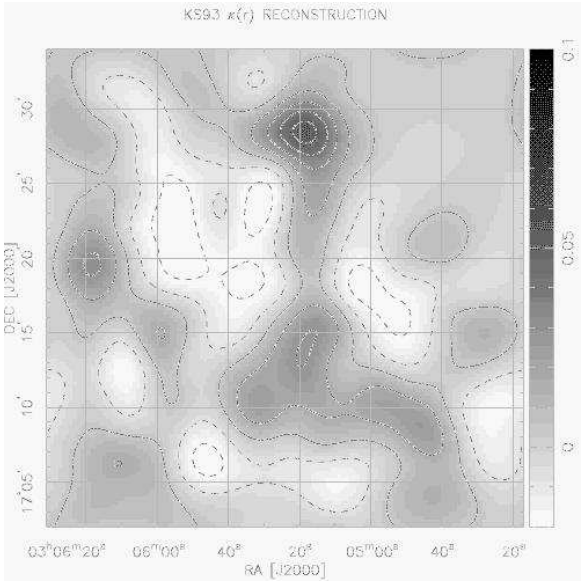

Fig. 11. Mass map of the supercluster MS0302+17, smoothed on a scale of 90 arcseconds (black indicates higher density). Each of the three densest blobs corresponds to a known galaxy cluster. From [74].

This result is robust: it remains when the source selection criteria are varied, and it disappears when the source positions are randomized. There appears to be a filament connecting two of the clusters, but the authors advise caution, as the signal-to-noise is low, and real filaments are not expected to have much contrast against all the other filaments and sheets expected to lie be-

tween sources and observer. An equally striking reconstruction of the Abell 901/902 supercluster was recently published [50]. The close correspondence of mass peaks and known clusters says something about the predictability of dark matter, as discussed below.

Despite these successes, it is worth remembering that a map (or radial profile) of κ is not a map of mass. κ can be converted to mass only with a careful calibration of the source redshifts, which must include an estimate of source contamination by the cluster itself. In massive clusters, magnification provides another source of error by increasing the mean redshift of sources which have been selected according to an apparent magnitude cut. This results in $\Sigma_{\rm crit}$ being a function of radius, as it decreases at small radii where the sources of a fixed magnitude tend to be more distant [47]. While much attention has been paid to optimizing the formal reconstruction methods, these more mundane problems require equal attention.

2.3 Mass and light

Does light trace mass ? The answer must be at least a qualified yes, because the projected shapes of cluster lenses tend to agree with the shapes suggested by their emitted light (Figure 12). However, the qualifications are important !

First, the correspondence holds only on scales larger than galaxies. The vast majority of visible-wavelength light from clusters comes from individual galaxies, not diffuse emission. Although weak lensing is not well suited to examine small scales, there is ample evidence from strong lensing that cluster mass distributions (like their X-ray emissions) do not peak on galaxy scales [123].

Second, not all light is equal. Blue light is dominated by very young stars, while established stellar populations which presumably trace mass better tend to be red. Hence small variations in star formation could scatter the ratio of mass to blue light M/L_B widely from system to system, but M/L_R should be much more stable. Unfortunately, the literature has a tradition of quoting M/L_B, which obfuscates the issue of whether mass is traced by light from established stellar populations. Compounding this confusion are different methods of computing rest-frame emission given the observed emission, and the occasional quotation of M/L at $z = 0$, meaning that L has first been adjusted to a value that it would have at $z = 0$ given passive evolution. Because stellar populations fade with time, L can decrease, and M/L increase, significantly from its *in situ* value.

Third, even if M/L were more consistently defined, it may not be constant spatially or as a function of scale. Although the literature is full of mass reconstructions which generally follow smoothed light distributions, the M/L found varies widely, from $\sim 80h$ to $\sim 800h$. Some of this is no doubt due to different methods mentioned above, but there is reason to believe that not all of it is. For example, the high value of $\sim 800h$ for MS1224+20 was

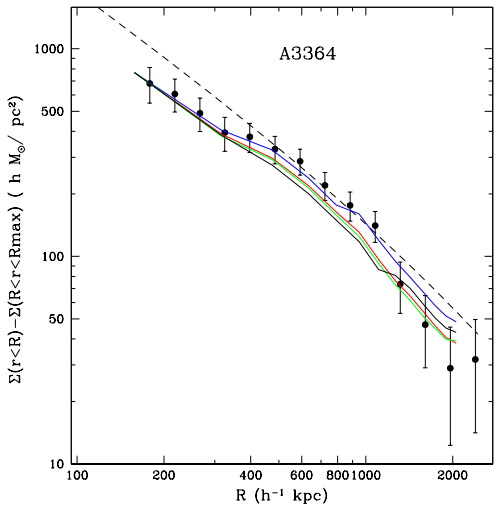

Fig. 12. Projected mass and light density profiles of Abell 3364. The light profiles were observed in observer-frame B_j (blue line), V (green), R (red), and I (black) filters, and computed in the same differential apertures used for the mass. The light profiles have each been shifted vertically to intersect the innermost mass point, hence they are in arbitrary units. Mass follows light surprisingly well on all measurable scales. The dotted line shows the shape of an isothermal profile, which is not quite a straight line with this estimator, to guide the eye (it has not been fit to the data). The two lowest mass points are approaching the level of systematic error estimated from the point-spread function. Note that in the aperture densitometry method, error bars on adjacent points are not independent, so that the errors should be thought of as a band. From [139].

found independently by two different investigators [45,46]. Also, attempts to uniformly treat samples of ~ 10 clusters have found a range of M/L within the samples [117,139]. There are some hints that some of the scatter may be due to a trend of increasing M/L with cluster mass.

If so, this follows a broader trend in which groups have lower M/L than the typical cluster [57], and typical galaxies have still lower M/L. The idea that M/L might depend on environment is called *bias*. Specifically, bias is when light is more concentrated than mass; the reverse, *antibias*, may also occur. However, indications of antibias can be understood as a simple consequence of stellar evolution and the choice of a blue bandpass [7]. In the end, M/L may be more a question of star formation history and bandpasses than of the nature of dark matter.

An alternative approach in terms of galaxy-mass correlations may offer more promise. Because large mass concentrations are clearly more associated

with early-type galaxies than with later types (the *morphology-density relation*), restricting the analysis to early-type galaxies might reveal a tighter relationship to mass. This was first done for the MS 0302 supercluster, shown in Figure 11 [74]. A cross-correlation between the projected mass density and that predicted from the early-type galaxies revealed a strong relationship, which did not vary with density, whereas a simple M/L would have acquired variations from the variations in star formation activity. This approach has been extended to the field, with similar conclusions [136]. Correlations between mass and light were found to be not so simple in the Abell 901/902 supercluster [50], but perhaps the difference is due to all light, not only that from early-type galaxies, being used in the Abell 901/902 work. Still, the correlation of mass with early-type galaxies must fail on some smaller scale, as we know that galaxy groups have mass but generally no early-type galaxies. Clearly more work is required in this area, as the correlation of mass with different types of emitters may provide clues to the nature of dark matter.

2.4 Clusters as cosmological probes

There is a hidden agenda behind all the effort that has gone into measuring cluster M/L: if cluster M/L is representative of the universe in general, the mean density of the universe Ω_m can be estimated simply by scaling the local luminosity density by this ratio. This is one version of the *fair sample hypothesis*, and it is one of the many ways to use clusters as cosmological probes. These can be divided roughly into methods that extrapolate from physical conditions in clusters (using clusters only because they facilitate certain measurements), and methods which use clusters to diagnose formation of structure in the universe.

Scaling the luminosity density by M/L is the simplest example of the first type of method. The fair sample hypothesis remains just a hypothesis, but it has nevertheless spawned many estimates of Ω_m, which tend to be $\lesssim 0.4$ [117,74,139,57]. However, the apparent variation of M/L with environment and age makes this approach suspect, and it is worth asking whether any property of clusters other than light could be used in a similar scaling argument.

The best candidate is the baryon fraction f_b, the ratio of baryonic mass to total mass. Because there is no reason to believe that infall into clusters favors baryons over dark matter or vice versa, f_b is plausibly equal to Ω_b/Ω_m. In addition to being plausible, the baryonic hypothesis is easier to investigate with simulations of structure formation, because tracking the baryons is easy in such simulations; the hard part is simulating star formation and the resultant light emission. With Ω_b fairly well known from Big Bang nucleosynthesis arguments [24], a determination of f_b would quickly yield Ω_m. Lensing can provide an estimate of the total mass, while SZE measurements can probe the dominant baryonic component, the intracluster gas. Simulations indicate that the combination should reveal f_b to 10% or better [142,39]. The first

results from real clusters (but with total mass estimated from X-ray emission rather than lensing) indicate $\Omega_m \sim 0.25$ [51]. It should be noted that any census of baryons is likely to be incomplete, as they can take many forms which are difficult to detect (brown dwarfs, planets, etc.). Hence this method provides a lower limit to f_b and an upper limit to Ω_m.

There is always a chance that physical properties of clusters such as f_b are simply not representative of the universe in general. A second and more powerful class of cosmological probe uses clusters as tracers of structure. Only their mass is important, and in particular, their *mass function*, the number density of clusters as a function of mass. The redshift evolution of the cluster mass function is a probe of Ω_m: all else being equal, a high-density universe should show more recent evolution than a low-density universe. In fact, it has been argued that the existence of even one massive cluster at high redshift (*e.g.* MS1054 at $z = 0.83$ [38] and now also ClJ1226.9+3332 at $z = 0.88$ [41]) demonstrates that evolution has not been as rapid as required if $\Omega_m = 1$ [19]. Conclusions based on a few massive clusters are suspect, however, because as the extreme tail of a distribution, their numbers are highly dependent on the assumption of Gaussianity in the primordial fluctuations. The argument can even be turned around: given an independent measure of Ω_m, cluster counts can put strong constraints on primordial non-Gaussianity [106,78]. With plentiful wide-field data now available and with weak lensing techniques having been honed on less massive clusters, it will soon be possible to construct an honest mass function, which will constrain both quantities [81]. The redshift evolution of the cluster mass function can also constrain dark energy [55,67]. Without any uniform weak-lensing cluster samples, though, we must defer this discussion to Future Prospects and turn our attention to progress in obtaining such a sample.

2.5 Shear-selected clusters

The use of clusters as cosmological probes centers on clusters as mass concentrations, not as collections of galaxies and gas. Yet all cluster samples compiled to date have been based on emitted light from galaxies (*e.g.* [2]) or from a hot intracluster medium [20]. Because these mechanisms do not involve dark matter, which is the dominant component by mass, a mass function based on these samples may well be biased. In addition, the r^{-2} falloff of emitted light implies that the high-redshift end of such samples will always be dominated by the most luminous clusters, which potentially introduces another bias. Shear-selected clusters are needed to investigate these potential biases and provide a clean mass function, and this is currently an active area in weak lensing.

First, a note on terminology. "Cluster" implies a collection of galaxies, but if a large mass concentration with no visible galaxies were to be identified, it would probably be called a "dark cluster". Although "dark matter halo" would be a more accurate term, it is used almost exclusively in theoretical

and computational papers, not observational work. Here we shall continue to use the term "cluster", but we emphasize that this is a working hypothesis. After large samples of shear-selected mass concentrations are thoroughly followed up with other methods, it will become clear if a different term is more appropriate.

Unfortunately, such samples are not available yet. Although many previously known clusters were studied with weak lensing in the 1990s, no surveys for new clusters were conducted, partly because of the small fields of view afforded by cameras on large telescopes until later in the decade, and partly because techniques needed to be proven on known clusters first. The first serendipitous detections of mass concentrations came when unexpected peaks appeared some distance from the target in mass reconstructions of known clusters.

In the first reported detection, a mass concentration was found projected near Abell 1942 ($7'$ from the center), and confirmed by a mass map constructed with data from a different camera and at a different wavelength [42]. There is no obvious concentration of galaxies associated with this mass, although the area does contain a poor group of galaxies and some weak X-ray emission. Because the redshift of the mass is unknown, its mass and M/L are also unknown. However, with an upper limit on the light in the area, the lower limit on M/L can be computed as a function of redshift. There are redshifts for which the object could have a reasonable M/L, around 400 [49]. This object is therefore not necessarily more dark than some X-ray selected clusters, which have M/L up to 600 or more [45,46] (see Section 2.3). If it is at the redshift of Abell 1942 ($z = 0.22$), its M/L is at least 600, which is very dark but just on the edge of the X-ray selected range, and perhaps explainable in a merger scenario with Abell 1942 proper.

In a second detection, Hubble Space Telescope (HST) imaging revealed an extra mass concentration one arcminute from the center of CL1604+4304 [128]. It is also seen in a second pointing shifted by $20''$, so it is likely to be real, but the interpretation is not clear. It seems likely to be substructure in the cluster rather than an independent structure, but the necessary followup is lacking. In a third case, serendipitous mass concentrations were found in a survey of known clusters [139]. Some of these corresponded to galaxy groups, which followup spectroscopy showed to be real and not associated with the target clusters, although in the same general redshift range (the range to which the lensing survey was of course most sensitive).

Although these hints were exciting, large "blank" fields (*i.e.* fields not selected to contain a known cluster) are more appropriate for finding unambiguously new clusters, and the first truly convincing shear-selected cluster was indeed discovered in such a field [138]. This object is clearly a cluster of galaxies (Figure 13), with a solid spectroscopic redshift (0.28), velocity dispersion (615 km s^{-1}), and lens redshift coinciding with the spectroscopic value (see Section 2.6). The M/L is at the high end of, but definitely within,

the range found for optically and X-ray selected clusters. Although clearly seen at visible wavelengths, no X-ray emission is detected at this position. A second shear-selected cluster has recently been assigned a spectroscopic redshift [137]. At $z = 0.68$, this cluster begins to fulfill the promise of lensing in terms of avoiding the r^{-2} falloff of methods which depend on emitted light.

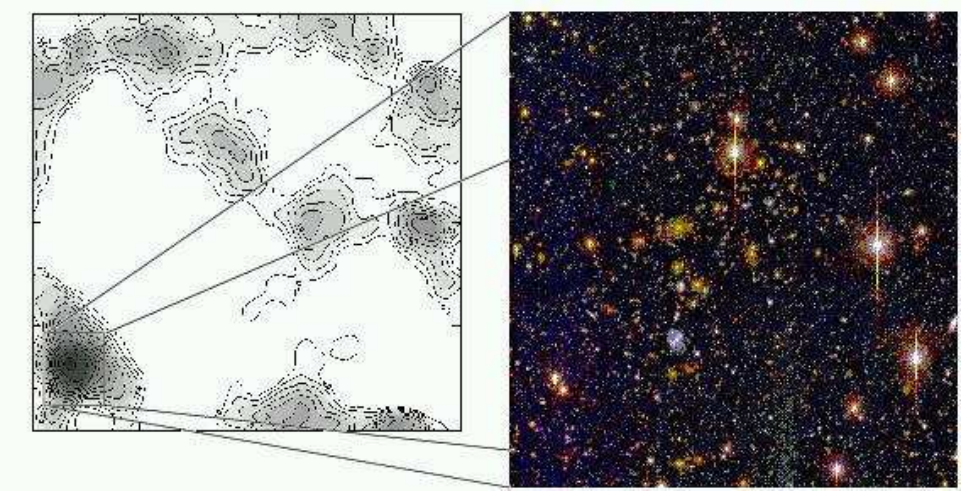

Fig. 13. A shear-selected cluster of galaxies. At left is a κ map of a $40'$ field not selected to contain a previously known cluster; black indicates higher density. The mass concentration at lower left corresponds to a cluster of galaxies (inset), spectroscopically confirmed at a redshift of 0.28 and a velocity dispersion of 615 km s^{-1}.

Finally, the most recent candidate makes perhaps the strongest case yet for a dark cluster [94]. A tangential alignment was found around a point in a randomly selected $50''$ STIS field, significant enough that the data allow only a 0.3% chance of this occuring randomly. There are indications of strong lensing as well. A nearby group of galaxies could provide enough mass to explain this only if its M/L is two orders of magnitude higher than expected. As with the first two cases cited above, more followup, including a lens redshift, is desperately needed to make sense of this candidate.

Thus far, the serendipitous shear detections present no clear pattern, apart from the feeling that these are not typical optically or X-ray selected clusters. There is no proof of truly dark clusters, something which only lensing could detect. Perhaps this is only for lack of followup. Yet, the detection of a truly dark cluster would reveal surprisingly little about the nature of dark matter. Rather, it would indicate that either baryons did not fall into the potential well created by the dark matter halo, or that star formation

failed there. These are intriguing scenarios, but they raise questions about baryons rather than answer questions about dark matter.

Meanwhile, there are several surveys of tens of square degrees currently underway [127,37], which will yield samples of dozens of shear-selected clusters, rather than a serendipitous few, and perhaps yield a better idea of typical and extreme shear-selected clusters. Much work remains in terms of settling on estimators which maximize detection of real mass concentrations while minimizing false positives. For example, do we simply look for peaks in convergence maps (or maps of some other quantity such as potential or aperture mass), or do we apply a matched filter, which implies that we know what we are looking for? While such work has been done theoretically and computationally [104], we must get our hands dirty with real samples before we can have much confidence in the scattered examples published as of today.

The advantages of shear selection in avoiding baryon and emitted-light bias are obvious, but no single cluster-finding technique will be completely unbiased. SZE selection [61] is an exciting new method which is also independent of emitted light. This is especially important in going to high redshift because of the r^{-2} falloff of emitted light. In this respect, SZE has the advantage because its background source is at a very high redshift (the cosmic microwave background at $z \sim 1100$), and because lensing is most efficient at detecting clusters at much lower redshift than the sources. However, lensing and SZE methods are so new that samples are not yet available. X-ray and optical selection are more established, and X-ray surveys have recently made great strides in detecting high-redshift clusters [20], indicating that it can compete with other methods at $z > 1$ despite the r^{-2} falloff of emitted light. X-ray emission has the additional advantage of depending on the square of the local density, making it less vulnerable to projection effects (although the density-squared dependence could be viewed as a disadvantage when trying to determine the total cluster mass). Table 1 summarizes the properties of these selection methods. In the end, comparison of differently selected samples will always be necessary, and much work remains to be done before we can claim that all the important biases are known.

Selection method	Projection effects?	Emitted light?	Baryon dependent?	Samples available now?
Optical	yes	yes	yes	yes
X-ray	no	yes	yes	yes
Lensing	yes	no	no	almost
SZE	yes	no	yes	almost

Table 1. Comparison of cluster selection methods.

2.6 Tomography with clusters

Judging from the examples of the previous section, followup and identification of shear-selected clusters will be more difficult than finding them. The most basic parameter, redshift, is unknown in several cases. Without a redshift, even the lens mass and M/L must remain unknown, leaving little solid information. A spectroscopic redshift is impossible in the case of the Abell 1942 field with no obvious lens-associated galaxies, and difficult in the CL1604+4304 field, with CL1604+4304 itself projected so nearby. Thus, there is a great need for a method of determining the lens redshift from the lensing information alone.

If sources can be differentiated by redshift, the redshift of a lens will be revealed by the way that shear increases with source redshift (Figure 3). Photometric redshifts are required for the sources, but this is straightforward if the deep imaging required for the shear measurement is extended to multiple filters. Two-filter imaging is routinely done anyway, to filter the sources based on color. Four filters is sufficient to provide photometric redshifts accurate to ~ 0.1 on each source, which is accurate enough given the large amount of shape noise on each galaxy and the breadth of the lensing kernel. This method has been demonstrated on one cluster [138] (Figure 14). The most likely lens redshift is within 0.03 of the spectroscopic redshift ($z = 0.28$), but the formal error estimate is ~ 0.1.

Obviously, lens redshifts cannot compete with cluster spectroscopic, or even photometric, redshifts with this level of precision. Some improvement is to be expected as photometric redshifts improve. For example, the filter set used was not designed to be optimal for photometric redshifts, but future large surveys will be paying close attention to this issue. Also, the work cited neglected to use sources which were undetected in one or more filters, but photometric redshifts are not impossible to assign to such sources, and their inclusion could improve the statistics. A rigorous treatment would take account of each source's photometric redshift error estimate, and so on. Work is needed on optimal lens redshift algorithms.

Still, lens redshifts can be useful even at this level of accuracy. The most basic use is to confirm the unstated assumption in all cluster weak lensing work to date—that the cluster *is* the lens, not merely in the same line of sight. While no doubt valid, confirmation of such basic assumptions is always welcome. Second, in cases where dark mass concentrations are found without any associated galaxies, a lens redshift is the only way to constrain the basic parameters of the mass concentration. Indeed, if there is any skepticism about such claims, it would be conclusively dispelled by demonstrating that the observed shear increases with source redshift in the predicted way. Third, large weak lensing surveys may find enough shear-selected clusters to make complete spectroscopic followup burdensome. In that case, rough lens redshifts may be good enough for examining statistics of many clusters, or at least for identifying the more interesting candidates for followup. There-

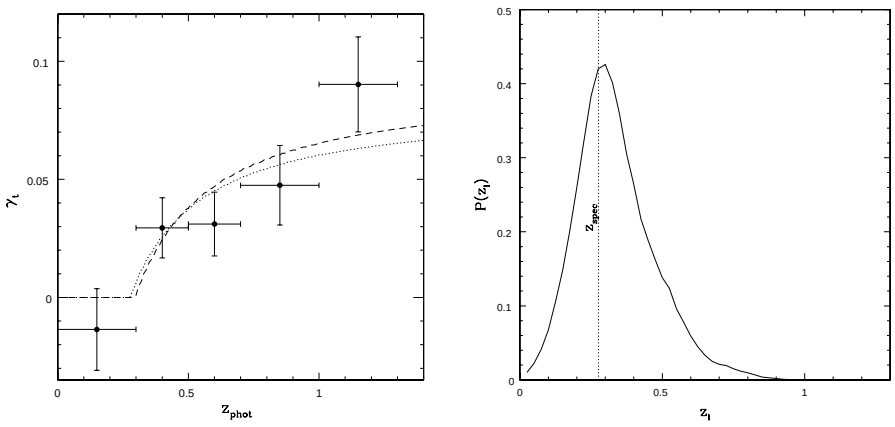

Fig. 14. *Left*: tangential shear as a function of source (photometric) redshift. The dotted line is the best fit for a lens at the cluster spectroscopic redshift of 0.28, while the dashed line is the best fit with the lens redshift derived from the lensing data alone. ($z = 0.30$). *Right*: The lens redshift probability distribution derived from the data at left. The method is promising: The most likely lens redshift is within 0.03 of the cluster spectroscopic redshift, but the width of the distribution is ~ 0.1, indicating the need for more precise data. From [138].

fore, this type of tomography will probably be a routine feature of future shear-selected surveys.

Finally, note how the spread in source redshifts would have caused more uncertainty in shear had source redshifts not been known in Figure 14. One way of improving cluster shear measurements, which until now have used at most a color cut to avoid contamination of the sources by cluster members, will be the use of photometric redshifts to weight sources. If Figure 14 is any guide, this might make an improvement of up to a factor of two.

3 Large-scale structure

Clusters are not the largest structures in the universe. Although it had long been known that clusters themselves tend to cluster, it was only in the 1980's that redshift surveys began to reveal apparently coherent structures—filaments and voids—on very large scales, up to ~ 50 Mpc. Current redshift surveys are producing impressive views of this foamy galaxy distribution out to ~ 600 Mpc, or $z \sim 0.2$ [100]. But what about the mass distribution ?

Simulations of cold dark matter show similar structures in mass (Figure 15). Furthermore, they show how the evolution of large-scale structure depends on cosmological parameters and on the nature of dark matter. Good measurements of large-scale structure evolution should therefore be able to

constrain cosmological parameters and the nature of dark matter through comparison with simulations. Weak lensing is a good candidate for such comparisons, because like the simulations it deals with mass, not galaxies; and because it can easily reach back to $z \sim 1-2$, providing a long baseline in cosmic time.

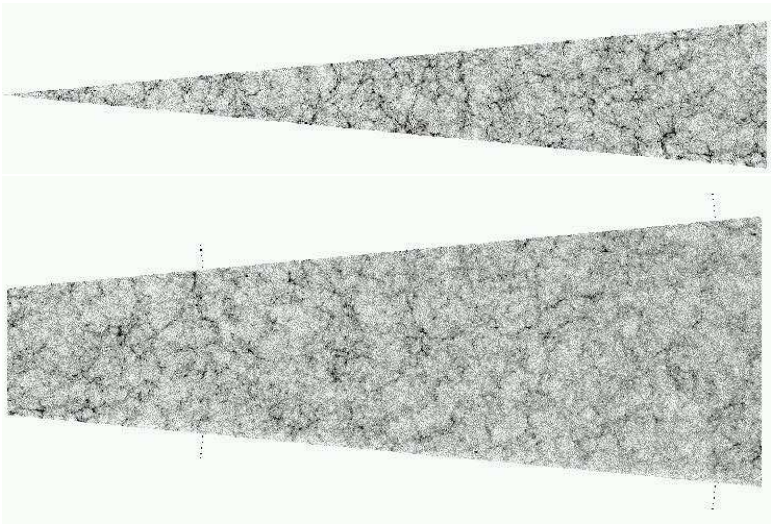

Fig. 15. Simulation from the Virgo collaboration showing the evolution of large-scale structure in a 7° slice of a Λ-dominated universe, with black indicating highest density. The dotted lines indicate $z = 1$ and $z = 2$. Adapted from [133].

Figure 15 illustrates just how many voids and filaments are expected to lie between us and a source at $z \sim 1$. Because of projection effects, weak lensing will never produce stunningly detailed three-dimensional mass maps to allow comparisons with such simulations. But weak lensing by large-scale structure does leave a statistical signature. This "cosmic shear" is a strong function of cosmological parameters, in particular, Ω_m and σ_8, the rms density variation on 8 Mpc scales, and thus is potentially a very useful cosmological tool. However, cosmic shear leaves a much weaker signal than do clusters, making detection more difficult and systematics more dangerous. The first detections of cosmic shear came in 2000, a decade after weak lensing by clusters was detected. For that reason, cosmic shear is just beginning to take its place in the cosmology toolbox.

3.1 Cosmic shear estimators

Cosmic shear, unlike the shear induced by clusters and groups, has no center, and that has led to the formulation of a variety of statistics different from

those used to analyze clusters. We summarize them here to provide the basis for interpreting the results presented here and in the literature. Each of the following statistics has advantages and disadvantages, and current practice is to report results in terms of several different estimators to verify robustness.

A few comments apply to all the estimators mentioned below. Current wide-field cameras have fields of view of $\sim 0.5°$, and all results to date have been reported on these or smaller scales. But a look at Figure 15, with its opening angle of 7°, shows that such small fields will give different results depending on where they happen to lie. Because of this *cosmic variance*, one such field cannot really constrain the cosmology. A sample of randomly chosen fields is required, with the field-to-field scatter in results giving some idea of the cosmic variance. Observed variance could also be due to problems with the instrument or telescope, so this is really an upper limit to the cosmic variance, but it is still a very useful number. Some groups are currently doing much larger fields by stitching together multiple pointings, sometimes with sparse sampling, but multiple, widely separated fields are still required to insure that cosmic variance has been beaten.

On small scales, the dominant statistical noise source is simply shape noise, but systematics are also larger here. Small-scale PSF variations cannot be mapped because the density of stars is too low; intrinsic alignments play a larger role on small scales; and comparison to theory (not necessarily simulations) is hampered by the difficulty of modeling the nonlinear collapse of dense regions. Results on scales $< 1'$ may say more about nonlinear collapse and possibly intrinsic alignments than about the cosmology.

Mean shear The mean shear in a field (simply averaging all sources) will in general be nonzero in the presence of lensing. However, it will tend to zero for a field of any significant size, so this statistic is of limited use. We mention it for completeness, as some early work with small fields of view used this statistic. But mean shear in a small field of view is difficult to interpret, as it could result from a single structure projected near the line of sight.

Shear variance The next logical step is to compute the variance, among a group of boxes of angular size ϕ, of the mean shear in each box. Because variance is a positive definite quantity, noise rectification must be subtracted off. The importance of the noise rectification term depends on the number of sources per box. Roughly speaking, for $\phi < 1'$, the noise rectification term is larger than the lensing signal itself, but for larger ϕ the lensing contribution dominates and at several arcminutes the noise correction becomes quite small. Thus, measurements of shear variance at $\phi < 1'$ should be treated with caution (in addition to the cautions cited above for small scales). Usually the results are presented as a function of ϕ, but the values for different ϕ have been computed from the same data. Hence the errors on the different scales are not independent, and results are sometimes less significant than

appears at first glance. The true significance of such results can be explored with bootstrap resampling. A rule of thumb suggested by bootstrap tests is that measurements on widely differing scales (factors of 10) are largely independent of each other even when computed from the same data.

Ellipticity correlations The observed ellipticities of lensed sources are correlated, so it is natural to construct a correlation function which measures this effect as a function of angular separation between sources. A simple correlation of the ellipticity components e_+ and $e_\times$ would have little physical meaning, though, as they depend on the orientation of the detector axes. If the components of a pair of galaxies are instead defined with respect to an imaginary line joining their centers, their correlation does have a physical significance [92]. In fact, three useful functions can be defined:

$$\begin{aligned} \xi_1(\theta) &\equiv \langle e_+^i e_+^j \rangle \\ \xi_2(\theta) &\equiv \langle e_\times^i e_\times^j \rangle \\ \xi_3(\theta) &\equiv \langle e_+^i e_\times^j \rangle \end{aligned} \tag{8}$$

where superscripts label the sources and brackets denote averaging over all pairs of galaxies $i \neq j$ with angular separation θ.

Like shear variance, lensing induces $\xi_1 > 0$ for all θ, but decreasing with θ. Unlike shear variance, the computation of ξ_1 does *not* result in a positive definite quantity, so spurious results may be easier to identify. The behavior of ξ_2 in the presence of lensing by large-scale structure is more interesting: at $\theta = 0$, ξ_2 and ξ_1 are equal, but ξ_2 drops more rapidly and goes negative at some θ which depends on the cosmology ($\sim 0.5 - 1^\circ$). ξ_3 is a control statistic. Unaffected by lensing, it should vanish in the absence of systematic errors or intrinsic alignments (this is equivalent to rotating one of each pair of galaxies by 45°). Taken together, these properties provide a signature with several lines of defense against systematic error.

Like shear variance, values for different θ are computed from the same data, so the same warnings about nonindependent angular bins apply. Unlike shear variance, though, there are two independent quantities (ξ_1 and ξ_2) at each θ, which can be checked against each other. For example, unless $\xi_1(0) = \xi_2(0)$ and $\xi_1(\theta) > \xi_2(\theta)$ for $\theta > 0$, the results are suspect. If these checks make sense, ξ_1 and ξ_2 can be combined into a single higher signal-to-noise measurement. In fact, shear variance can be understood as $\xi_1(0)+\xi_2(0)$, convolved with a square window function of width ϕ.

Aperture mass The aperture mass statistic M_{ap} was designed to address the nonlocality of shear. It is a generalization of the ζ statistic already mentioned in the context of clusters [45,73]. As in the ζ statistic, tangential shear is computed in a circular aperture, but here it is weighted with a compensated filter function; the weight is positive at the center of the aperture and negative

at the edges, for a total weight of zero. This has the remarkable property of making adjacent apertures nearly independent, whereas the shear in adjacent apertures is highly correlated. M_{ap} was first suggested as a way of looking for clusters in wide-field images [112] , and later its variance $\langle M_{ap}^2 \rangle$ was proposed as a measure for cosmic shear [113] (note that because the total weight vanishes, so does the expectation value: $\langle M_{ap} \rangle = 0$). Although aperture mass tends to be noisier than the other estimators, its compensating virtue is that measurements on different scales are almost completely independent.

Other estimators All of the above estimators are at most two-point statistics. Higher-order statistics have been proposed. For example, values of the projected mass field (κ or M_{ap}) should have a skewness due to many somewhat underdense regions (voids) and a few extreme overdense regions. This skewness depends on Ω_m and the matter power spectrum in a different way than does the variance, leading to suggestions that together they could constrain both quantities [14]. This non-Gaussianity may also be revealed through morphological analysis of convergence fields [110,121,52]. All this remains largely theoretical, as these high-order statistics are in practice noisier than the two-point estimators which are providing the first detections of cosmic shear. There is a recent claim of detection of this non-Gaussian signature [13], but an accurate measurement of skewness requires that rare massive halos be present in the sample [32], hence a very large area is required. The best way to the power spectrum itself may be through maximum likelihood fitting to the shear data [66,31].

3.2 Observational status

Although the idea of weak lensing by large-scale structure was first suggested in the 1960's [80,53], the effect escaped detection for over three decades. The first attempts at detection gave null results [79,130], which is not surprising given the subtleness of the effect ($\sim 1\%$ shear) and the lack of sensitivity and nonlinearity of photographic plates. The first analysis of CCD data, albeit with the narrow field afforded by CCDs in 1994, also yielded only upper limits [97]. With the advent of large-format CCD mosaics, detection was inevitable, and four groups [140,131,6,75] announced detections in the span of one month in 2000.

Their results are summarized in terms of shear variance in Figure 16. The groups used four different cameras on three different telescopes, with different observed bandpasses and data reduction procedures and analysis techniques, yet the results were in good agreement. This has been taken as proof that instrumental effects and systematic errors have been vanquished, but in fact, the results should *not* agree if the data and source selection resulted in different mean source redshifts. The fact that the areal density of sources used was similar for all four groups suggests that the source redshifts

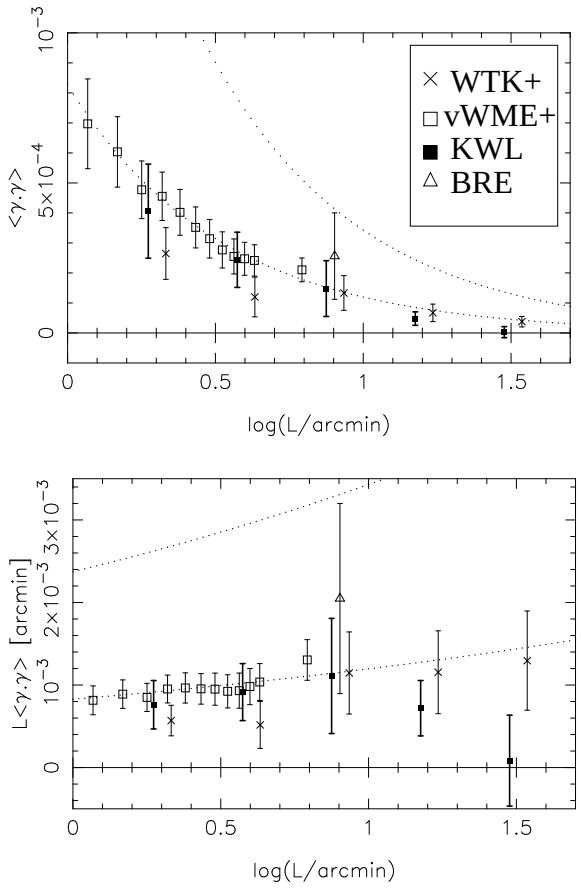

Fig. 16. First detections of cosmic shear, in terms of shear variance versus angular scale. The results of four different groups using three different telescopes and four different cameras are shown, with good agreement. Note that different angular bins from the same experiment are not independent. The dotted lines are for two different source redshift distributions (lower, $\langle z \rangle = 1$; upper, $\langle z \rangle = 2$) in a ΛCDM universe. Adapted from [75].

were similar despite the different approaches. But the possibility remains that different source redshifts are hiding some disagreement.

Nevertheless, all results point to a low-Ω_m universe. Figure 16 shows the good fit to ΛCDM. It is difficult to constrain Λ with these measurements, but shear variance should scale roughly with Ω_m, so it is clear that $\Omega_m = 1$, for example, is ruled out. While this was no surprise, it signaled the emergence of cosmic shear as a new way to constrain Ω_m, completely independent of traditional methods (supernovae, CMB, age of the oldest stars in conjunction with the Hubble constant, etc.).

Since then, there have been further detections both in ground-based [87,58] and space-based data [105]. The state of the art is a many-sigma detection (whatever estimator is chosen) over 6.5 deg^2 leading to quantitative constraints in the Ω_m, σ_8 plane [132]. There is a significant degeneracy between Ω_m and σ_8; the first generation of cosmic shear papers simply assumed a value of σ_8 consistent with the local abundance of clusters. At the same time, efforts to improve the signal-to-noise of cosmic shear measurements by decomposition of cosmic shear into E and B modes are underway and appear to have met with success [103].

Currently, several large (tens of square degrees) surveys are under way with the goal of very high signal-to-noise analyses of cosmic shear [127,37,58]. At the same time, the question of how accurate the measurements can ultimately get is being explored [5,43,82]. However, a word of caution is in order. Such analyses tend to ignore the fact that the source redshift distribution is not well known. The putative accuracy of current and near-future cosmic shear measurements thus tends to be far too optimistic. For the moment, the most accurate measurements in an absolute sense will be those which are no deeper than current redshift surveys. Within a few years, though, this will probably change as photometric redshifts are used to estimate the source redshift distribution accurately enough. Photometric redshifts, in fact, will enable probing the redshift evolution of cosmic shear; division of sources into just two or three redshift bins can greatly improve the measurements of cosmological parameters, specifically Ω_Λ by a factor of ~ 7 [62].

4 Future prospects

4.1 New applications

It is impossible to predict what new applications weak lensing might find, but it is worth discussing one example of an interesting new direction: constraints on theories of gravity. It is unlikely that weak lensing will serendipitously reveal some new feature of gravity, because the lenses through which we look are not well calibrated. But given an alternative theory of gravity, we can ask if weak lensing observations are consistent with other observations.

Modified gravity is an attempt to explain differences between light distributions and inferred mass distributions without invoking dark matter. It is possible to modify Newtonian gravity to account for some of the observed differences such as flat rotation curves in galaxies, but a general correlation between mass and light remains. If lensing were to find severe discrepancies between mass and light, such as a dark cluster or clear misalignment of cluster mass and light axes, this would represent a serious blow to modified gravity [115]. There are some promising dark cluster candidates (Section 2.5), and Abell 901b presents mass and light axes which apparently differ [50], but there are no bulletproof examples. Weak lensing surveys of significant areas are only now underway, so it will take some time before dark clusters can

be ruled in or out with much confidence. Note that dark clusters are not *expected* in dark matter scenarios; mass concentrations should accumulate enough baryons to become visible, if only in X-rays. Thus an absence of dark clusters would not favor modified gravity over dark matter, but their presence of would disprove modified gravity as currently envisioned.

Recently, the first quantitative predictions of weak lensing in modified gravity scenarios were published. Modified gravity, by increasing the strength of gravity on large scales, would greatly enhance cosmic shear, inconsistent with measurements. Thus, at the large scales probed by cosmic shear, the r^{-2} force law cannot be modified—if gravity does depart from r^{-2}, it is only on scales from 10 h^{-1} kpc to 1 h^{-1} Mpc [134]. Weak (and strong) lensing can also address modified gravity by constraining halo flattening [96]. Weak lensing by large-scale structure can also provide a test of higher-dimensional gravity [129].

4.2 New instruments

Today's surveys of tens of square degrees will take years to find perhaps dozens of shear-selected clusters and put some constraints on w, the dark energy equation of state, as well as Ω_m [55]. Tight constraints on both would require a very deep survey of 1000 deg^2 [67], taking decades with current telescopes and instruments. The latest generation of 8-m class telescopes does not really help, as their fields of view are small, typically $\sim 10'$ or less (an exception is the Subaru telescope which has a $\sim 24'$ field of view with its SuPrime camera, a likely source of weak-lensing results in the near future). A new generation of wide-field telescopes specifically designed for surveys will dramatically accelerate our ability to do cosmology with large lensing surveys. These surveys will cover an area comparable to that of SDSS, but much more deeply. The first of these to be funded is VISTA, a 4-m telescope with a 1° field of view currently in the design stage, but apparently it will be infrared-only, limiting its usefulness for weak lensing. LSST, an 8-m class telescope with a 3° field of view [126] and concentrating on visible wavelengths, may also be built within a decade.

Figure 17 shows the potential of a 1000 deg^2 survey which LSST could easily accomplish. Of course, predictions such as these depend on the extrapolation of $\sqrt{n}$ statistics to extremely large areas, so it is wise to ask what systematic effects might provide a higher noise floor. Early work on cosmology constraints from cluster counts assumed NFW profiles for all clusters [81]. It was then realized that the profile makes a big difference, so that cluster counts may tell us more about dark-matter profiles than about cosmology [9]. However, new estimators have been proposed to circumvent this problem [55]. Careful attention must also be paid to the issue of completeness versus false positives in cluster-detection surveys [135]. Still, by the time LSST starts operation, these issues may be worked out, and it may be well to survey all

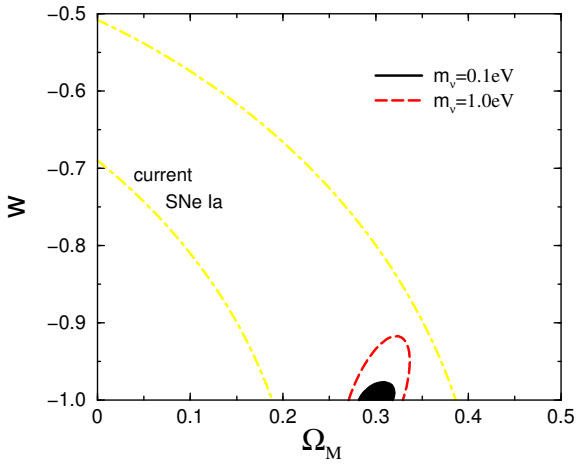

Fig. 17. 68% confidence limit constraints on Ω_M and w for two values of m_ν, for a weak lensing survey of 1000 deg^2 down to $R = 27$, with photometric redshifts providing the source redshift distribution. Current 1-σ constraints from type Ia supernovae are shown for comparison. From [67].

the sky visible from the site. Such a survey would also provide a shear power spectrum comparable in accuracy to the CMB power spectra of today.

Another probe of cosmology which may become feasible with such massive surveys involves the angular power spectrum of clusters. The linear part of this power spectrum is essentially a standard ruler calibrated by the CMB, so that a power spectrum of clusters at a particular redshift yields the angular diameter distance to that redshift. A very large survey (~ 4000 deg^2) could determine this as a function of redshift, which of course would yield an absolute calibration of the distance scale and the Hubble constant [33].

4.3 New algorithms

The combination of lensing data with other types of data has been an active theoretical area recently, and some of these algorithms will soon prove themselves observationally. Lensing plus SZE measurements of clusters will reveal the baryon fraction in that environment, perhaps leading to a new estimate of Ω_m from baryon scaling arguments—or perhaps leading to new aspects of cluster formation. Combinations of lensing and SZE plus X-ray data will help deproject cluster mass and gas distributions [142].

Cross-correlation of lensing by large-scale structure with the CMB will reveal parameters largely hidden from traditional CMB analyses, such as dark energy, the end of the dark ages, and the gravitational wave amplitude [64,65]. Lensing of the CMB itself may be detectable by the Planck satellite

and constrain the amplitude of mass fluctuations between us and $z \sim 1000$ [120], but this may have to wait for even higher-sensitivity CMB probes [63].

References

1. H. Abdelsalam, P. Saha, L. Williams: AJ **116**, 1541 (1998)
2. G. Abell, H. Corwin, R. Olowin: ApJS **70**, 1 (1989)
3. S. Allen: MNRAS **296**, 392 (1998)
4. S. W. Allen, R. W. Schmidt, A. C. Fabian: MNRAS **328**, L37 (2001)
5. D. Bacon, A. Refregier, D. Clowe, R. Ellis: MNRAS **325**, 1065 (2001)
6. D. Bacon, A. Refregier, R. Ellis: MNRAS **318**, 625 (2000)
7. N. Bahcall, R. Cen, R. Davé, J. Ostriker, Q. Yu: ApJ **541**, 1 (2000)
8. C. Balland, A. Blanchard: ApJ **487**, 33 (1997)
9. M. Bartelmann, L. King, P. Schneider: A&A **378**, 361 (2001)
10. M. Bartelmann: A&A **303**, 643 (1995)
11. M. Bartelmann, R. Narayan: ApJ **451**, 60 (1995)
12. M. Bartelmann, R. Narayan, S. Seitz, P. Schneider: ApJ **473**, 610 (1996)
13. F. Bernardeau, Y. Mellier, L. van Waerbeke: ApJL, submitted, astro-ph/0201032 (2002)
14. F. Bernardeau, L. van Waerbeke, Y. Mellier: A&A **322**, 1 (1997)
15. G. Bernstein, M. Jarvis: AJ **123**, 583 (2002)
16. A. Biviano: to appear in "Tracing Cosmic Evolution with Galaxy Clusters", ASP Conference Series; astro-ph/0110053 (2001)
17. J. Blakeslee: astro-ph/0108253 (2001)
18. R. Blandford, A. Saust, T. Brainerd, J. Villumsen: MNRAS **251**, 600 (1991)
19. P. Bode, N. Bahcall, E. Ford, J. Ostriker: ApJ **551**, 15 (2001)
20. S. Borgani, L. Guzzo: Nature **409**, 39 (2001)
21. T. Brainerd, C. Wright, D. Goldberg, J. Villumsen: ApJ **524**, 9 (1999)
22. T. Broadhurst, A. Taylor, J. Peacock: ApJ **438**, 49 (1995)
23. M. Brown, A. Taylor, N. Hambly, S. Dye: MNRAS, submitted, astro-ph/0009499 (2000)
24. S. Burles, K. Nollett, M. Turner: ApJL **552**, L1 (2001)
25. P. Catelan, M. Kamionkowski, R. Blandford: MNRAS, **320L**, 7 (2001)
26. R. Cen: ApJ **485**, 39 (1997)
27. T.-C. Chang, A. Refregier: ApJ, **570**, 447 (2002)
28. D. Clowe, P. Schneider: A&A **379**, 384 (2001)
29. J. Cohen, D. Hogg, R. Blandford, L. Cowie, E. Hu, A. Songaila, P. Shopbell, K. Richberg: ApJ **538**, 29 (2000)
30. A. Connolly, I. Csabai, A. Szalay, D. Koo, R. Kron, J. Munn: AJ **110**, 2655 (1995)
31. A. Cooray, W. Hu: ApJ **554**, 56 (2001)
32. A. Cooray, W. Hu: ApJ **548**, 7 (2001)
33. A. Cooray, W. Hu, D. Huterer, M. Joffre: ApJL **557**, 7 (2001)
34. R. Crittenden, P. Natarajan, U. Pen, T. Theuns: ApJ **559**, 552 (2001)
35. R. Crittenden, P. Natarajan, U. Pen, T. Theuns: ApJ **568**, 20 (2002)
36. R. Croft, C. Metzler: ApJ **545**, 561 (2000)
37. http://terapix.iap.fr/Descart

38. M. Donahue, G. Voit, I. Gioia, G. Luppino, J. Hughes, J. Stocke: ApJ **502**, 550 (1998)
39. O. Dore, F. Bouchet, Y. Mellier, R. Teyssierm: A&A **375**, 14 (2001)
40. S. Dye, A.N. Taylor, T.R. Greve, O.E. Rognvaldsson, E. van Kampen, P. Jakobsson, V.S. Sigmundsson, E.H. Gudmundsson, J. Hjorth: A&A **386**, 12 (2002)
41. H. Ebeling, L.R. Jones, B.W. Fairley, E. Perlman, C. Scharf, D. Horner: ApJL **548**, 23 (2001)
42. T. Erben, L. van Waerbeke, Y. Mellier, P. Schneider, J.-C. Cuillandre, F.J. Castander, M. Dantel-Fort: A&A **355**, 23 (2000)
43. T. Erben, L. van Waerbeke, E. Bertin, Y. Mellier, P. Schneider: A&A **366**, 717 (2001)
44. E. Falco, M. Gorenstein, I. Shapiro: ApJ **289**, L1 (1985)
45. Fahlman, G., Kaiser, N., Squires, G., & Woods, D.: ApJ **437**, 56 (1994)
46. P. Fischer: AJ **117**, 2024 (1999)
47. P. Fischer, J. A. Tyson: AJ **114**, 14 (1997)
48. D. Goldberg, M. Natarajan: ApJ, submitted, astro-ph/0107187 (2001)
49. M. Gray, R. Ellis, J. Lewis, R. McMahon, A. Firth: MNRAS **325**, 111 (2001)
50. M. Gray, A. Taylor, K. Meisenheimer, S. Dye, C. Wolf, E. Thommes: ApJ **568**, 141 (2002)
51. L. Grego, J. Carlstrom, E. Reese, G. Holder, W. Holzapfel, M. Joy, J. Mohr, S. Patel: ApJ **552**, 2 (2001)
52. A. Guimaraes: MNRAS, submitted, astro-ph/0202507 (2002)
53. J. Gunn: ApJ **147**, 61 (1967)
54. A. Heavens, A. Refregier, C. Heymans: MNRAS **319**, 649 (2000)
55. J. Hennawi, V. Narayanan, D. Spergel, I. Dell'Antonio, V. Margoniner, J. A. Tyson, D. Wittman: BAAS **199**, 1608 (2001)
56. H. Hoekstra: A&A **370**, 743 (2001)
57. H. Hoekstra, M. Franx, K. Kuijken, R.G. Carlberg, H.K.C. Yee, H. Lin, S.L. Morris, P.B. Hall, D.R. Patton, M. Sawicki, G.D. Wirth: ApJL **548**, L5 (2001)
58. H. Hoekstra, H. Yee, M. Gladders: to appear in the proceedings of the STScI 2001 spring symposium "Dark Universe", astro-ph/0106388 (2001)
59. D. Hogg: astro-ph/9905116 (1999)
60. D. Hogg, *et al.*: AJ **115**, 1418 (1998)
61. G. Holder, J. Mohr, J. Carlstrom, A. Evrard, E. Leitch: ApJ **544**, 629 (2000)
62. W. Hu: ApJL **522**, L21 (1999)
63. W. Hu: PRD **62**, 3007 (2000)
64. W. Hu: ApJL **557**, L79 (2001)
65. W. Hu: Phys. Rev. D, submitted, astro-ph/0108090 (2001)
66. W. Hu, M. White: ApJ **554**, 67 (2001)
67. D. Huterer: ApJ, submitted, astro-ph/0106399 (2001)
68. N. Kaiser: ApJ **388**, 272 (1992)
69. N. Kaiser: ApJ **498**, 26 (1998)
70. N. Kaiser: in *Gravitational Lensing: Recent Progress and Future Goals*, ASP Conference Proceedings, Vol. 237. Edited by Tereasa G. Brainerd and Christopher S. Kochanek. San Francisco: Astronomical Society of the Pacific, ISBN: 1-58381-074-9, p.269 (2001)
71. N. Kaiser: ApJ **537**, 555 (2000)
72. N. Kaiser, G. Squires: ApJ **404**, 441 (1993)

73. N. Kaiser, G. Squires, T. Broadhurst: ApJ **449**, 460 (1995)
74. N. Kaiser, G. Wilson, G. Luppino, L. Kofman, I. Gioia, M. Metzger, H. Dahle: ApJ, submitted, astro-ph/9809268 (1998)
75. N. Kaiser, G. Wilson, G. Luppino: ApJL, submitted, astro-ph/0003338 (2000)
76. C. Keeton: astro-ph/0102341 (2001)
77. L. King, P. Schneider: A&A **369**, 1 (2001)
78. A. Knebe, R. Islam, J. Silk: MNRAS **326**, 109 (2001)
79. J. Kristian: ApJ **147**, 864 (1967)
80. J. Kristian, R. Sachs: ApJ **143**, 379 (1966)
81. G. Kruse, P. Schneider: MNRAS **302**, 821 (1999)
82. K. Kuijken: astro-ph/0007368 (2000)
83. J. Lee, U. Pen: ApJ **555**, 106 (2000)
84. J. Lee, U. Pen: ApJL **567**, L111 (2000)
85. R. Lynds, V. Petrosian: BAAS **18**, 1014 (1986)
86. J. Mackey, M. White, M. Kamionkowski: MNRAS, **332**, 788 (2002)
87. R. Maoli, L. van Waerbeke, Y. Mellier, P. Schneider, B. Jain, F. Bernardeau, T. Erben, B. Fort: A&A **368**, 766 (2001)
88. C. Mayen, G. Soucail: A&A **361**, 415 (2000)
89. Y. Mellier: ARAA **37**, 127 (1999)
90. C. Metzler, M. White, C. Loken: ApJ **547**, 560 (2001)
91. J. Miralda-Escudé: ApJ **370**, 1 (1991)
92. J. Miralda-Escudé: ApJ **380**, 1 (1991)
93. J. Miralda-Escudé, A. Babul: ApJ **449**, 18 (1995)
94. J.-M. Miralles, T. Erben, H. Haemmerle, P. Schneider, R.A.E. Fosbury, W. Freudling, N. Pirzkal, B. Jain, S.D.M. White: A&A submitted, astro-ph/0202122 (2002)
95. R. Moessner, B. Jain: MNRAS **294**, 291 (1998)
96. D. J. Mortlock, E. L. Turner: MNRAS **327**, 552 (2001)
97. J. Mould, R. Blandford, J. Villumsen, T. Brainerd, I. Smail, T. Small, W. Kells: MNRAS **271**, 31 (1994)
98. J. Navarro, C. Frenk, S. White: MNRAS **275**, 720 (1995)
99. J. Navarro, C. Frenk, S. White: ApJ **490**, 493 (1997)
100. J. Peacock *et al.*: Nature **410**, 169 (2001)
101. P. J. E. Peebles: *Principles of Physical Cosmology*, Princeton University Press (1993)
102. U. Pen, Li, U. Seljak: ApJL **543**, L107 (2000)
103. U. Pen, L. van Waerbeke, Y. Mellier: ApJ **567**, 31 (2002)
104. K. Reblinsky, M. Bartelmann: A&A **345**, 1 (1999)
105. J. Rhodes, A. Refregier, E. Groth: ApJL **552**, 85 (2001)
106. J. Robinson, E. Gawiser, J. Silk: ApJ **532**, 1 (2000)
107. A. Refregier: MNRAS, submitted, astro-ph/0105178 (2001)
108. A. Refregier, D. Bacon: MNRAS, submitted, astro-ph/0105179 (2001)
109. W. Saslaw: *Gravitational Physics of Stellar and Galactic Systems*, Cambridge University Press (1987)
110. J. Sato, M. Takada, Y. P. Jing, T. Futamase: ApJL **551**, L5 (2001)
111. P. Schneider, S. Seitz: A&A **294**, 411 (1995)
112. P. Schneider: MNRAS **283**, 837 (1996)
113. P. Schneider, L. van Waerbeke, B. Jain, G. Kruse: MNRAS **296**, 873 (1998)
114. P. Schuecker, H. Bohringer, H. Reiprich, L. Feretti: A&A **378**, 408 (2001)

115. J. Sellwood, A. Kosowsky: to appear in *The Dynamics, Structure & History of Galaxies*, G. S. Da Costa & E. M. Sadler, eds, ASP Conference Series; astro-ph/0109555 (2002)
116. E. Sheldon, *et al.*: ApJ **554**, 881 (2001)
117. I. Smail, R. Ellis, A. Dressler, W. Couch, A. Oemler, R. Sharples, H. Butcher: ApJ **479**, 70 (1997)
118. G. Soucail, B. Fort, Y. Mellier, J. Picat: A& A **172**, L14 (1987)
119. G. Squires, N. Kaiser: ApJ **473**, 65 (1996)
120. M. Takada, T. Futamase: ApJ **546**, 620 (2001)
121. A. Taruya, M. Takada, T. Hamana, I. Kayo, T. Futamase: ApJ, submitted, astro-ph/0202090 (2002)
122. J. A. Tyson: AJ **96**, 1 (1988)
123. J. A. Tyson, G. Kochanski, I. Dell'Antonio: ApJ **498**, 107 (1998)
124. J. A. Tyson, P. Seitzer: ApJ **335**, 552 (1988)
125. J. A. Tyson, R. Wenk, F. Valdes: ApJL **349**, L1 (1990)
126. J. A. Tyson, D. Wittman, J. R. P. Angel: to appear in proceedings of the Dark Matter 2000 conference (Santa Monica, February 2000) to be published by Springer, astro-ph/0005381 (2000)
127. J. A. Tyson, D. Wittman, I. Dell'Antonio, A. Becker, V. Margoniner, DLS Team: BAAS **199**, 10113 (2001)
128. K. Umetsu, T. Futamase: ApJL **539**, L5 (2000)
129. J.-P. Uzan, F. Bernardeau: Phys. Rev. D 64, 3004 (2001)
130. F. Valdes, J. Jarvis, J. A. Tyson: ApJ **271**, 431 (1983)
131. L. van Waerbeke *et al.*: A&A **358**, 30 (2000)
132. L. van Waerbeke *et al.*: A&A **374**, 757 (2001)
133. http://www.mpa-garching.mpg.de/Virgo/virgoproject.html
134. M. White, C. S. Kochanek: ApJ **560**, 539 (2001)
135. M. White, L. van Waerbeke, J. Mackey: astro-ph/0111490 (2001)
136. G. Wilson, N. Kaiser, G. Luppino: ApJ **556**, 601 (2001)
137. D. Wittman, V. Margoniner, J. A. Tyson, J. Cohen, I. Dell'Antonio: ApJL in prep (2002)
138. D. Wittman, J. A. Tyson, V. Margoniner, J. Cohen, I. Dell'Antonio: ApJL **557**, L89 (2001)
139. D. Wittman, I. Dell'Antonio, J. A. Tyson, G. Bernstein, P. Fischer, D. Smith: in *Constructing the Universe with Clusters of Galaxies*, IAP 2000 meeting, Paris, France, July 2000, Florence Durret & Daniel Gerbal (Eds.) (2000)
140. D. Wittman, J. A. Tyson, D. Kirkman, I. Dell'Antonio, G. Bernstein: Nature **405**, 143 (2000)
141. D. York *et al.*: AJ **120**, 1579 (2000)
142. S. Zaroubi, G. Squires, G. de Gasperis, A. Evrard, Y. Hoffman, J. Silk: ApJ **561**, 600 (2001)