A compressed sensing approach to 3D weak lensing

A. Leonard, F.-X. Dupé, and J.-L. Starck

Laboratoire AIM, UMR CEA-CNRS-Paris 7, Irfu, SAp/SEDI, Service d'Astrophysique, CEA Saclay, 91191 Gif-sur-Yvette Cedex, France
e-mail: adrienne.leonard@cea.fr; francois-xavier.dupe@lif.univ-mrs.fr

Received 6 July 2011 / Accepted 29 November 2011

ABSTRACT

Context. Weak gravitational lensing is an ideal probe of the dark universe. In recent years, several linear methods have been developed to reconstruct the density distribution in the Universe in three dimensions, making use of photometric redshift information to determine the radial distribution of lensed sources.

Aims. We aim to address three key problems seen in these methods; namely, the bias in the redshifts of detected objects, the line-of-sight smearing seen in reconstructions, and the damping of the amplitude of the reconstruction relative to the underlying density. We also aim to detect structures at higher redshifts than have previously been achieved, and to improve the line-of-sight resolution of our reconstructions.

Methods. We considered the problem under the framework of compressed sensing (CS). Under the assumption that the data are sparse or compressible in an appropriate dictionary, we constructed a robust estimator and employ state-of-the-art convex optimisation methods to reconstruct the density contrast. For simplicity in implementation, and as a proof of concept of our method, we reduced the problem to one dimension, considering the reconstruction along each line of sight independently. We also assumed an idealised survey in which the redshifts of sources are known.

Results. Despite the loss of information inherent in our one-dimensional implementation, we demonstrate that our method is able to accurately reproduce cluster haloes up to a redshift of $z_{cl} = 1.0$, deeper than state-of-the-art linear methods. We directly compare our method with these linear methods, and demonstrate minimal radial smearing and redshift bias in our reconstructions, as well as a reduced damping of the reconstruction amplitude as compared to the linear methods. In addition, the CS framework allows us to consider an underdetermined inverse problem, thereby allowing us to reconstruct the density contrast at finer resolution than the input data.

Conclusions. The CS approach allows us to recover the density distribution more accurately than current state-of-the-art linear methods. Specifically, it addresses three key problem areas inherent in linear methods. Moreover, we are able to achieve super-resolution and increased high-redshift sensitivity in our reconstructions.

Key words. gravitational lensing: weak – methods: statistical – techniques: image processing – cosmology: observations – galaxies: clusters: general – large-scale structure of Universe

1. Introduction

Weak gravitational lensing has become a powerful tool for studying the dark universe, allowing us to place constraints on key cosmological parameters, and offering the possibility to place independent constraints on the dark energy equation of state parameter, w (Levy & Brustein 2009; Hoekstra & Jain 2008; Munshi et al. 2008; Albrecht et al. 2006; Peacock et al. 2006; Schneider 2006; Van Waerbeke & Mellier 2003).

Until recently, weak lensing studies considered the shear signal, and recovered the mass distribution, in two-dimensional projection (see Schneider 2006, for a review of weak lensing). However, with improved data quality and wide-band photometry, it is now possible to recover the mass distribution in three dimensions by using photometric redshift information to deproject the lensing signal along the line of sight (Simon et al. 2009; Massey et al. 2007a,b; Taylor et al. 2004).

Under the assumption of Gaussian noise, various linear methods have been developed to recover the 3D matter distribution, which rely on the construction of a pseudo-inverse operator to act on the data, and which include a penalty function encoding the prior that is to be placed on the signal (VanderPlas et al. 2011; Simon et al. 2009; Castro et al. 2005; Taylor et al. 2004; Bacon & Taylor 2003; Hu & Keeton 2002).

These methods produce promising results; however, they show a number of problematic artefacts. Notably, structures detected using these methods are strongly smeared along the line of sight, the detected amplitude of the density contrast is damped (in some cases, very strongly), and the detected objects are shifted along the line of sight relative to their true positions. These effects result from the choice of method used; Simon et al. (2009) note that their choice of filter naturally gives rise to a biased solution, and VanderPlas et al. (2011) suggest that linear methods might be fundamentally limited in the resolution attainable along the line of sight as a result of the smearing effect seen in these methods.

Furthermore, these methods are restricted to deal solely with the overdetermined inverse problem. In other words, the resolution obtainable on the reconstruction of the density is limited to be, at best, equal to that of the input data. Thus, the resolution of the reconstruction is entirely limited by the quality of the data and its associated noise levels, with no scope for improvement by judicious choice of inversion or denoising method.

In this paper, we consider the deprojection of the lensing signal along the line of sight as an instance of compressed sensing, where the sensing operator models the line-of-sight integration of the matter density giving rise to the lensing signal. For

simplicity in implementation, and as a proof of concept of our method, we considered only the line-of-sight transformation, reducing the 3D weak lensing problem to a one-dimensional inversion, with each line of sight in an image treated as independent. We note, however, that the algorithm presented here can be easily generalised to three dimensions.

We adopted a sparse prior on our reconstruction, and used a state-of-the-art iterative reconstruction algorithm drawn from convex analysis, optimisation methods and harmonic analysis set within a compressed sensing framework. This enables us to find a robust estimator of the solution without requiring any direct prior knowledge of the statistical distribution of the signal. We note that the compressive sensing framework allows us to consider an underdetermined inverse problem, whereby we are able to obtain higher resolution on our reconstructions than that provided by the input data.

This method produces reconstructions with minimal bias and smearing in redshift space, and with reconstruction amplitudes $\sim 75\%$ of the true amplitude (or better, in some cases). This is a significant improvement over current linear methods, despite the adoption of a simplified, one-dimensional algorithm. In addition, our method exhibits an apparent increased sensitivity to high-redshift structures, as compared with linear methods. Our reconstructions do exhibit some noise, with false detections appearing along a number of lines of sight. However, these tend to be localised to one or two pixels, rather than coherent structures, and we expect improved noise control with a full 3D implementation.

We note that we do not include photometric redshift errors in our simulations and, consequently, the simulations shown here should be considered to be idealised. Ma et al. (2006) have presented a method to account for these errors in lensing measurements. This method has been used by various authors in studies on real data (see, e.g. Simon et al. 2011), and is straightforward to implement in the algorithm presented here. A full treatment of photometric redshift errors, which is essential for the method to be useful on real data, will be presented in an upcoming work.

This paper is structured as follows. In Sect. 2, we outline the weak lensing formalism in three dimensions, and outline several linear inversion methods to solve the 3D weak lensing problem. We introduce our compressed sensing framework and describe our proposed algorithm in Sect. 3. In Sect. 4, we discuss practical considerations in implementing the method, and describe our simulated dataset. In Sect. 5 we demonstrate the performance of our algorithm in reconstructing simulated cluster haloes at various redshifts. We conclude with a discussion of our results and future applications in Sect. 6.

Throughout the text, we assumed a Λ CDM cosmology, with $\Omega_\Lambda = 0.736$, $\Omega_M = 0.264$, $h = 0.71$, $\sigma_8 = 0.801$, consistent with the WMAP-7 results (Larson et al. 2011).

2. 3D weak lensing

The distortion of galaxy images caused by the weak lensing effect is described, on a given source plane, by the Jacobian matrix of the coordinate mapping between source and image planes:

$$A = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix}, \quad (1)$$

where κ is the projected dimensionless surface density, and $\gamma = \gamma_1 + i\gamma_2$ is the complex shear. The shear is related to the convergence via a convolution in two dimensions:

$$\gamma(\theta) = \frac{1}{\pi} \int d^2\theta' \mathcal{D}(\theta - \theta') \kappa(\theta'), \quad (2)$$

where

$$\mathcal{D}(\theta) = \frac{1}{(\theta^*)^2}, \quad (3)$$

$\theta = \theta_1 + i\theta_2$, and an asterisk $*$ represents complex conjugation.

The convergence, in turn, can be related to the 3D density contrast $\delta(\mathbf{r}) \equiv \rho(\mathbf{r})/\bar{\rho} - 1$ by

$$\kappa(\theta, w) = \frac{3H_0^2\Omega_M}{2c^2} \int_0^w dw' \frac{f_K(w')f_K(w-w')}{f_K(w)} \frac{\delta[f_K(w')\theta, w']}{a(w')}, \quad (4)$$

where H_0 is the Hubble parameter, Ω_M is the matter density parameter, c is the speed of light, $a(w)$ is the scale parameter evaluated at comoving distance w , and

$$f_K(w) = \begin{cases} K^{-1/2} \sin(K^{1/2}w), & K > 0 \\ w, & K = 0 \\ (-K)^{-1/2} \sinh([K]^{1/2}w) & K < 0, \end{cases} \quad (5)$$

gives the comoving angular diameter distance as a function of the comoving distance and the curvature, K , of the Universe.

If the shear (or convergence) data is divided into N_{sp} redshift bins, and the density contrast reconstruction is divided into N_{lp} redshift bins (where N_{sp} is not necessarily equal to N_{lp}), we can write the convergence $\kappa^{(i)}$ on each source plane as

$$\kappa^{(i)}(\theta) \simeq \sum_{\ell=1}^{N_{lp}} Q_{i\ell} \delta^{(\ell)}(\theta), \quad (6)$$

where

$$Q_{i\ell} = \frac{3H_0^2\Omega_M}{2c^2} \int_{w_\ell}^{w_{\ell+1}} dw \frac{\bar{W}^{(i)}(w)f_K(w)}{a(w)}, \quad (7)$$

and

$$\bar{W}^{(i)}(w) = \int_0^{w^{(i)}} dw' \frac{f_K(w-w')}{f_K(w')} \left(p(z) \frac{dz}{dw} \right)_{z=z(w')}. \quad (8)$$

Thus, for each line of sight, Eq. (6) describes a matrix multiplication, encoding a convolution along the line of sight. It is the inversion of this transformation:

$$\kappa(z) = \mathbf{Q}\delta(z), \quad (9)$$

that is the focus of this paper. We note that the inversion of Eq. (2) can be straightforwardly performed on each source plane in Fourier space.

2.1. Linear inversion methods

We focus here on the methods presented in Simon et al. (2009) and VanderPlas et al. (2011). For a review of other linear methods, the reader is referred to Hu & Keeton (2002).

The 3D lensing problem is effectively one of observing the density contrast convolved with the linear operator $\mathbf{R}$, and contaminated by noise, which is assumed to be Gaussian. Formally, we can write

$$\mathbf{d} = \mathbf{R}\mathbf{s} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2), \quad (10)$$

where $\mathbf{d}$ is the observation, $\mathbf{s}$ the real density and $\boldsymbol{\varepsilon}$ the Gaussian noise.

The general idea behind linear inversion methods is to find a linear operator $\mathbf{H}$ that acts on the data vector to yield a solution that minimises some functional, such as the variance of the residual between the estimated signal and the true signal, subject to some regularisation or prior-based constraints.

The simplest instance of such a linear operation is an inverse variance filter (Aitken 1934), which weights the data only by the noise covariance, and places no priors on the signal itself:

$$\hat{\mathbf{s}}_{IV} = [\mathbf{R}^\dagger \mathbf{\Sigma}^{-1} \mathbf{R}]^{-1} \mathbf{R}^\dagger \mathbf{\Sigma}^{-1} \mathbf{d}, \quad (11)$$

where $\mathbf{\Sigma} \equiv \langle \mathbf{nn}^\dagger \rangle$ gives the covariance matrix of the noise.

This method proves problematic when the matrices involved are non-invertible, such as when there are degeneracies inherent in the allowed solution. In order to make the problem invertible, some regularisation must be introduced. Simon et al. (2009) opt to use a Saskatoon filter (Tegmark 1997; Tegmark et al. 1997), which combines a Wiener filter and an inverse variance filter, with a tuning parameter α introduced that allows switching between the two. This gives rise to a minimum variance filter, expressed as

$$\hat{\mathbf{s}}_{MV} = [\alpha \mathbf{1} + \mathbf{S} \mathbf{R}^\dagger \mathbf{\Sigma}^{-1} \mathbf{R}]^{-1} \mathbf{S} \mathbf{R}^\dagger \mathbf{\Sigma}^{-1} \mathbf{d}, \quad (12)$$

where $\mathbf{S} \equiv \langle \mathbf{ss}^\dagger \rangle$ encodes prior information about the signal covariance, and $\mathbf{1}$ is the identity matrix.

VanderPlas et al. (2011) have recently proposed a filter based on the singular value decomposition (SVD) of the inverse variance filter of Eq. (11). Under this formalism, we can write

$$\hat{\mathbf{s}}_{IV} = \mathbf{V} \mathbf{\Lambda}^{-1} \mathbf{U}^\dagger \mathbf{\Sigma}^{-1/2} \mathbf{d}, \quad (13)$$

where we have decomposed the matrix $\tilde{\mathbf{R}} \equiv \mathbf{\Sigma}^{-1/2} \mathbf{R} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^\dagger$, $\mathbf{U}^\dagger \mathbf{U} = \mathbf{V}^\dagger \mathbf{V} = \mathbf{1}$ and $\mathbf{\Lambda}$ is the square diagonal matrix of singular values $\lambda_i = \Lambda_{ii}$.

Their filtering consists in defining a cutoff value σ_{cut} , and truncating the matrices to remove all singular values with $\sigma_i < \sigma_{\text{cut}}$. This effectively reduces the noise by removing the low singular values, which translate into high values in the inversion.

VanderPlas et al. note that the SVD decomposition is computationally intensive, and while they do describe a method to speed up the process, it may not be practical to use this method on large images.

Similar considerations must be made when using any of the linear methods described above, because these all involve matrix inversion, which is an $O(N^3)$ process. While optimisations can be found, these methods become excessively time- and computer-intensive when large datasets are considered. This, in effect, limits the resolution attainable using these methods.

Moreover, as discussed extensively in Simon et al. (2009) and VanderPlas et al. (2011), these linear methods give rise to a significant bias in the location of detected peaks, damping of the peak signal, and a substantial smearing of the density along the line of sight. The compressed sensing (CS) theory, described below, allows us to address the lensing inversion problem under a new perspective, and we will show that these three aspects are significantly improved using a non-linear CS approach.

3. Compressed sensing approach

Linear methods are easy to use, and the variance of each estimator is quite direct to compute. Furthermore, these methods generally rely on very common tools with efficient implementation. However, they are not the most powerful, and including non-Gaussian priors is difficult, especially when these priors

imply non-linear terms. Obviously, using better-adapted priors is required for building a more robust estimator. We adopted a compressed sensing approach to construct an estimator that exploits the sparsity of the signal that we aim to reconstruct. The estimator is modelled as an optimisation problem that is solved using recent developments from convex analysis and splitting methods.

3.1. Compressed sensing theory

We considered some data Y_i ($i \in [1, \dots, m]$) acquired through the linear system

$$Y = \mathbf{\Theta} X, \quad (14)$$

where $\mathbf{\Theta}$ is an $m \times n$ matrix. Compressed sensing (Candès & Tao 2006; Donoho 2006) is a sampling/compression theory based on the sparsity of the observed signal, which shows that, under certain conditions, one can exactly recover a k -sparse signal (a signal for which only k pixels have values different from zero, out of n total pixels, where $k < n$) from $m < n$ measurements.

This recovery is possible from undersampled data only if the sensing matrix $\mathbf{\Theta}$ verifies the *restricted isometry property* (RIP) (see Candès & Tao 2006, for more details). This property has the effect that each measurement Y_i contains some information about the entire pixels of X ; in other words, the sensing operator $\mathbf{\Theta}$ acts to spread the information contained in X across many measurements Y_i .

Under these two constraints – sparsity and a transformation meeting the RIP criterion – a signal can be recovered exactly even if the number of measurements m is much smaller than the number of unknown n . This means that, using CS methods, we will be able to outperform the well-known Shannon sampling criterion by far.

The solution X of (14) is obtained by minimizing

$$\min_X \|\mathbf{X}\|_1 \quad \text{s.t.} \quad Y = \mathbf{\Theta} X, \quad (15)$$

where the ℓ_1 norm is defined by $\|\mathbf{X}\|_1 = \sum_i |X_i|$. The ℓ_1 norm is well-known to be a sparsity-promoting function; i.e. minimisation of the ℓ_1 norm yields the most sparse solution to the inverse problem. Many optimisation methods have been proposed in recent years to minimise this equation. More details about CS and ℓ_1 minimisation algorithms can be found in Starck et al. (2010).

In real life, signals are generally not “strictly” sparse, but are *compressible*; i.e. we can represent the signal in a basis or frame (Fourier, Wavelets, Curvelets, etc.) in which the curve obtained by plotting the obtained coefficients, sorted by their decreasing absolute values, exhibits a polynomial decay. Note that most natural signals and images are compressible in an appropriate basis.

We can therefore reformulate the CS equation above (Eq. (15)) to include the data transformation matrix $\mathbf{\Phi}$:

$$\min_\alpha \|\alpha\|_1 \quad \text{s.t.} \quad Y = \mathbf{\Theta} \mathbf{\Phi} \alpha, \quad (16)$$

where $X = \mathbf{\Phi}^* \alpha$, and α are the coefficients of the transformed solution X in $\mathbf{\Phi}$, which is generally referred to as the *dictionary*. Each column represents a vector (also called an *atom*), which ideally should be chosen to match the features contained in X . If $\mathbf{\Phi}$ admits a fast implicit transform (e.g. Fourier transform, Wavelet transform), fast algorithms exist to minimise Eq. (16).

One problem we face when considering CS in a given application is that very few matrices meet the RIP criterion. However, it has been shown that accurate recovery can be obtained as long the mutual coherence between $\mathbf{\Theta}$ and $\mathbf{\Phi}$,

$\mu_{\Theta, \Phi} = \max_{i,k} |\langle \Theta_i, \Phi_k, \rangle|$, is low (Candes & Plan 2010). The mutual coherence measures the degree of similarity between the sparsifying basis and the sensing operator. Hence, in its relaxed definition, we consider a linear inverse problem $Y = \Theta \Phi X$ as being an instance of CS when

1. the problem is underdetermined;
2. the signal is compressible in a given dictionary Φ ;
3. the mutual coherence $\mu_{\Theta, \Phi}$ is low. This will happen every time the matrix $\mathbf{A} = \Theta \Phi$ has the effect of spreading out the coefficients α_j of the sparse signal on all measurements Y_i .

Most CS applications described in the literature are based on such a soft CS definition. Compressed sensing was introduced for the first time in astronomy for data compression (Bobin et al. 2008; Barbey et al. 2011), and a direct link between CS and radio-interferometric image reconstruction was recently established in Wiaux et al. (2009), leading to dramatic improvement thanks to the sparse ℓ_1 recovery (Li et al. 2011).

3.1.1. CS and weak lensing

The 3D weak lensing reconstruction problem can be seen to completely meet the soft-CS criteria above. Indeed,

1. the problem is undetermined because we seek a higher resolution than initially provided by the observations, which are noise-limited;
2. the matter density is seen to be sparsely distributed, showing clusters connected by filaments surrounding large voids;
3. the lensing operator encoded by the matrix $\mathbf{Q}$ in Eq. (9) (or equivalently, the combination $\mathbf{P}_{\gamma\kappa} \mathbf{Q}$, where $\mathbf{P}_{\gamma\kappa}$ encodes the convolution in Eq. (2)) spreads out the information about the underlying density in a compressed sensing way.

To highlight point (3) above, Fig. 1 shows the rows (top panel) and columns (bottom panel) of the transformation matrix, $\mathbf{Q}$, encoding weak lensing along the line of sight. The top panel shows the lensing efficiency kernel, which reflects the sensitivity of a given source plane to the shearing effects of lenses at various redshifts. This is the broad convolution kernel of Eq. (4). The bottom panel shows the effect of a discrete lens at a given redshift on source planes, and demonstrates that localised lenses give rise to effects that are non-local and affect all sources at redshifts greater than the lens redshift.

3.1.2. Sparsity prior

Sparse priors have been shown to be very useful in regularising ill-posed inverse problems (see Fadili & Starck 2009, and references therein). In addition, a sparse prior using a wavelet basis has been used in many areas of signal processing in astronomy, such as denoising, deconvolution, and inpainting to recover missing data in images (Starck et al. 2010). The idea underlying these priors is that there exists a dictionary in which a given dataset is sparsely represented. The dictionary used should therefore match the shapes of the structures that we aim to detect as closely as possible.

Many experiments, such as N -body simulations, have shown that the matter density in the universe is mainly distributed in localised clusters, which are connected by thin filaments. Because structures in the Universe appear to be physically sparse, we may therefore assume that the matter density is sparse in a domain adapted for cluster- and curve-like structures. These domains exist and can be constructed by gathering several well-chosen

Fig. 1. *Top panel:* lensing efficiency kernel for given source planes as a function of lens redshift. *Bottom panel:* lensing efficiency of a given lens plane as a function of source redshift.

transforms inside a dictionary (e.g. a combination of wavelets and curvelets). The dictionary should be chosen to match the type of structure we aim to recover as closely as possible (e.g. isotropic structures for wavelets, filaments for curvelets), and may be adapted to include structures that are not sparse in the direct domain, but which may be sparsely represented in an appropriate dictionary.

Regularisation using a sparsity constraint can be understood through a Bayesian framework. We assume that the distribution of the solution in the sparsifying dictionary has a Laplacian distribution. This is equivalent to constraining the ℓ_1 norm, which promotes sparsity. A Gaussian assumption, in contrast, constrains the ℓ_2 norm, and leads to the standard Wiener filtering, a Tikhonov solution or an SVD solution, depending on the way we constrain the solution.

For specific classes of inverse problem, it has even been shown that the sparse recovery leads to the exact solution of the problem (compressed sensing). This behaviour does not exist with any other prior than sparsity. Obviously, the sparse recovery will be optimal when the signal is sparse, in the same way that the Wiener filter is optimal when the signal (and noise) is Gaussian. Because the Laplacian assumption (in an appropriate space such as wavelets) is more applicable than the Gaussian distribution, restoration of astronomical data is generally much more efficient using sparsity.

This explains why wavelets have been so successful for astronomical image restoration/detection. For the reconstruction of clusters along the line of sight, we are in a perfect situation for sparse recovery because clusters are not resolved owing to

the bin size, and they can therefore be modelled as Dirac δ -functions. We therefore take Φ to be a δ -function dictionary. Clearly, in this case, the pixel domain is especially appropriate for sparse recovery. We believe that the sparse prior is a much better model for this kind of data compared to previous methods with implicit Gaussian assumptions. Our results appear to support this claim.

We note that the method presented here is somewhat similar to the point-source reconstruction method described in [Hu & Keeton \(2002\)](#), though they used an ℓ_2 minimisation combined with a strict prior on the number of haloes present along a line of sight. Our method, in contrast, places no priors on the number of structures along the line of sight.

3.2. Problem statement

Under the CS framework, the reconstruction of the matter density amounts to finding the most sparse solution that is consistent with the data. There are many different ways to formulate such an optimisation problem, and we opt for the following:

$$\min_{s \in \mathbb{R}^n} \|\Phi^* s\|_1 \quad \text{s.t.} \quad \frac{1}{2} \|d - Rs\|_{\Sigma^{-1}}^2 \leq \epsilon \quad s \in C, \quad (17)$$

where the term to minimise is a sparsity-penalty function over the dictionary coefficients. The second term in Eq. (17) above is a data fidelity constraint, with Σ being the covariance matrix of the noise and ϵ the allowed distance between the estimation and the observation, while the final term forces the solution to have values inside a given interval, usually $C = [-1, +\infty]^n$ for matter overdensity.

Note that this latter constraint, encoding a hard minimum on the signal to be recovered, is not possible with linear methods, and is therefore an additional strength of our method. Enforcing these physical constraints on our solution helps to ensure the recovery of the most physically compelling solution given the data.

In the compressed sensing literature, we can find equivalent writing of the problem (17), but with a focus on the data fidelity. We preferred to directly seek for the sparsest solution for a given freedom (i.e. ϵ) on the distance to the observation, because this distance is usually easier to tune than an equilibrium between a regularization and the data fidelity terms.

Note this optimisation problem can be equivalently expressed in a Bayesian framework, assuming Gaussian noise and a Laplacian distribution of the dictionary coefficients.

In order to solve (17), we used the primal-dual splitting method of [Chambolle & Pock \(2011\)](#), which is described in full in Appendix A. The algorithm is iterative, and effectively splits the problem into two parts, applying the two constraints in Eq. (17) separately.

On each iteration, the estimate of the reconstruction is compared with the data, and the data fidelity constraint (the second constraint in Eq. (17)) is applied by projection of the residual onto an ℓ_2 ball of radius ϵ . Independently, the estimate of the solution is projected onto the sparsifying basis, and sparsity is imposed through soft-thresholding, which minimises the ℓ_1 norm of the basis coefficients. The threshold level is set by the parameter λ , which aids in controlling the noise. The noise modelling is discussed in Sect. 4.3, whilst the algorithm, and all practical considerations related to its implementation, are discussed in detail in Appendix B. Figure 2 shows a simplified schematic of the algorithm used.

Fig. 2. Simplified schematic of the reconstruction algorithm described in the text.

4. Implementation of the algorithm

We first consider the simulated data to be used in the remainder of this paper, before discussing some practical considerations important when implementing the algorithm.

4.1. Cluster simulations

In order to test our method, we needed to simulate a realistic data set. To this end, we consider here a fiducial survey with a background galaxy number density $n_g = 100 \text{ arcmin}^{-2}$ distributed in redshift according to

$$p(z) \propto z^2 e^{-(1.4z/z_0)^{1.5}}, \quad (18)$$

with $z_0 = 1.0$ ([Taylor et al. 2007](#); [Kitching et al. 2011](#)), and the distribution truncated at a maximum redshift of $z_{\text{max}} = 2$. Figure 3 shows this probability distribution, normalised arbitrarily.

We took the intrinsic dispersion in shear measurements to be $\sigma_\gamma = 0.2$, and considered a field of $1^\circ \times 1^\circ$ divided into a grid of 60×60 pixels. These parameters were chosen to mimic the data quality expected from next-generation surveys such as Euclid ([Refregier et al. 2010](#)) and LSST ([Ivezic et al. 2008](#); [LSST Science Collaboration et al. 2009](#)).

In each simulated image, one or more clusters were generated following an NFW density profile with $M_{200} = 10^{15} M_\odot$, $c = 3$ binned into N_{sp} redshift bins. The effective convergence and shear were computed by integrating the lensing signal within each source redshift bin, and Gaussian noise was added, scaled appropriately by the number density of galaxies within that bin.

Fig. 3. Redshift probability distribution $p(z)$ of the sources as a function of source redshift z for the simulations described in the text.

4.2. Reconstructions in 1D

As noted in Sect. 2, the 3D weak lensing problem can be reduced to a one-dimensional problem, by taking as our data vector the (noisy) lensing convergence along each line of sight, which is related to the density contrast through Eq. (9). Therefore, we took $\mathbf{d} = \kappa_{ij}(z)$ and $\mathbf{R} = \mathbf{Q}$, and considered each line of sight in our images independently. Furthermore, as discussed previously, we took Φ to be a δ -function dictionary.

In our simulations, clusters were placed into a region where the mean density in the absence of the cluster is equal to the mean density of the Universe at that redshift. In other words, δ is constrained to be greater than zero in all our simulations. Therefore, the projection onto the convex set $\mathcal{C}$ in algorithm 2 applies a positivity constraint at each iteration.

Clearly, a one-dimensional implementation throws away information, because we do not account at all for the correlation between neighbouring lines of sight that will arise in the presence of a large structure in the image; however, reducing the problem to a single dimension is fast and easy to implement, and allows us to test the efficacy of the algorithm using a particularly simple basis function through which we impose sparsity. A fully 3D treatment of the problem, with more accurate noise modelling (see below) will be the subject of a future work.

However, the algorithm used is entirely general; therefore, with appropriate choice of a 3D basis set and taking $\mathbf{d} = \gamma(\theta, z)$ and $\mathbf{R} = \mathbf{P}_{\gamma}\mathbf{Q}$, one can implement this algorithm as a fully 3D treatment of the data with no modification to the algorithm itself.

4.3. Noise modelling and control

Noise model for the data

We assumed that the redshifts of the sources are known exactly, so there is no correlation between the noise in each source bin. Therefore, the covariance matrix of the noise along the line of sight is diagonal, with

$$\Sigma_{ii} = \sigma_g^2(z_i) = \frac{\sigma_\gamma^2}{n_g(z_i)A_{\text{pix}}}, \quad (19)$$

where A_{pix} is the pixel area, $n_g(z_i)$ is the number density of sources in the bin at redshift z_i , and σ_γ is the intrinsic dispersion in galaxy ellipticity, taken throughout to be 0.2.

Fig. 4. Reconstructions for two lines of sight obtained with varying ϵ . The *top* row shows the input data and reconstructed convergence vector, while the *bottom* row shows the reconstructed density contrast.

In the algorithm presented above, the covariance matrix is used in the evaluation of the data fidelity constraint. Note that the covariance matrix is only diagonal if the galaxy redshifts are known exactly. In practice, photometric redshift errors mean that each redshift slice in the data is likely to be contaminated with a few galaxies whose redshift error bars overlap with neighbouring redshift bins. In this case, the covariance matrix will have additional, non-diagonal elements that are non-zero. This is straightforward to model, however, for the chosen method of photometric redshift estimations, and our algorithm is entirely general regarding the form of the covariance matrix. Therefore, the problem of photometric redshift errors is readily tractable in our method, and will be presented in a future work.

Noise control in the algorithm

The noise in the reconstruction is controlled and suppressed by two parameters in the algorithm described in Fig. 2 and Appendix B. The first, and most important of these parameters is the data fidelity control parameter, ϵ .

This parameter controls how well the data are fitted by the reconstruction, with $\epsilon = 0$ implying a perfect fit to the data, which is not possible in the presence of noise. Figure 4 demonstrates the effect of varying ϵ in the reconstruction of two lines of sight from our simulations.

Clearly, when ϵ is small, the algorithm attempts to fit each data point more closely which, in the presence of noise, can result in overfitting of the data (as seen in line of sight 1) and hence false detections along the line of sight. On the other hand, a large ϵ may result in a solution that is a poor fit to the data (as seen in line of sight 2).

The second parameter used to control the noise is the soft threshold parameter λ , which is used in the algorithm to impose the sparsity prior. A threshold set excessively high will result in a null solution, whilst a threshold set fairly low will allow for more false detections of noise peaks along a given line of sight.

The appropriate value for this threshold should be related to the expected fluctuations in the density contrast resulting from noise variations, and should scale with the signal-to-noise ratio in the image.

Note that while ϵ strongly affects the accuracy of the estimation in reproducing the underlying density contrast, λ simply affects the sparsity of the solution. In other words, changing λ will not greatly affect the reconstructions of true density peaks, but may affect the number of false detections and noise peaks seen. Also note that a thresholding λ does not imply that density peaks with $\delta < \lambda$ will not be detected, because soft thresholding is only applied to one part of the estimate of the solution.

5. Results

5.1. Comparison with linear methods

Firstly, to demonstrate the effectiveness of our method, we compared our method directly with the linear methods of Simon et al. (2009) and VanderPlas et al. (2011). Because these linear methods are only defined for $N_{\text{lp}} \leq N_{\text{sp}}$, we considered the case of $N_{\text{lp}} = N_{\text{sp}} = 20$, and generated a single cluster halo at a redshift of $z_{\text{cl}} = 0.25$ following an NFW halo profile with $M_{200} = 10^{15} M_{\odot}$ and $c = 3$.

In the SVD method of VanderPlas et al. (2011), we take $v_{\text{cut}} = 1 - \sum_{i=1}^n \sigma_i^2 / \sum_{i=1}^{n_{\text{max}}} \sigma_i^2 = 0.01$, and in the transverse and radial Wiener filtering methods of Simon et al. (2009), we take the tuning parameter to be $\alpha = 0.05$ in both cases. For our CS approach, we take the soft threshold parameter $\lambda = 8$, and the data fidelity control parameter $\epsilon = 3$.

Note that while the linear methods take $\mathbf{d} = \gamma(\theta, z)$, $\mathbf{R} = \mathbf{P}_{\gamma\kappa}\mathbf{Q}$, our method takes $\mathbf{d} = \kappa(\theta, z)$, $\mathbf{R} = \mathbf{Q}$ as before. The noise levels in each case are identical.

The results are presented in Figs. 5 and 6. Figure 5 presents the 2D projections of the reconstructions, computed by integrating the reconstruction along each line of sight, and the 1D reconstructions along the four central lines of sight. In the 3D renderings of Fig. 6, the reconstructions from our method, the SVD method, and the radial Wiener filter method are thresholded at $\delta = 3$ (i.e. the plot only shows $\delta_{\text{rec}} \geq 3$), and each is smoothed with a Gaussian of width $\sigma = 0.7$ pix in all three directions. The reconstruction from the transverse Wiener filter method is heavily damped with respect to the amplitude of the density contrast; a threshold of $\delta = 5 \times 10^{-6}$ is chosen in this case, and no smoothing is applied because the reconstruction already shows a very smooth distribution.

The SVD method appears, in the 1D plots, to identify the correct redshift of the cluster, with a small amount of line-of-sight smearing, but the plots show a prominent high-redshift peak along the line of sight. We note that it may be possible to remove this false detection by raising v_{cut} , but at the cost of increased line-of-sight smearing (see VanderPlas et al. 2011). The two-dimensional projection consisting of the integrated signal along each line of sight is seen to be more noisy than our results. The 3D rendering shows that the SVD method does well at identifying and localising the cluster, but the resulting reconstruction does suffer from widespread noise at a moderate level.

The radial and tangential Wiener filter methods show very little noise in the 2D projections, and the radial Wiener filter shows very little smearing of the reconstruction along the line of sight. However, neither Wiener method recovers correctly the redshift of the cluster. While the transverse Wiener filter shows very little noise, it exhibits a broad smearing in both the radial and transverse directions, and a heavy damping of the

Fig. 5. Comparison of our method with the linear methods as labelled. The *left column* shows the 2D projection of the reconstruction, while the *right column* shows the 1D reconstructions along the four central lines of sight (dashed lines). Note that, owing to the amplitude damping effect in SVD and Wiener reconstructions, the *y-axis* scaling is different in each of the line-of-sight plots.

amplitude of the reconstruction. The results obtained using the radial Wiener filter are considerably better, with very little noise seen in the reconstruction, and with the cluster seen to be well-localised in the radial direction. However, the amplitude of the cluster reconstruction is a factor of ~ 2.5 smaller than the input density.

Our results are seen to suffer from several prominent, pixel-scale false detections along noisy lines of sight not associated with the cluster. However, along the four central lines of sight

Fig. 6. 3D rendering of the reconstruction of a $z_{\text{cl}} = 0.25$ cluster using our method (*top*), the SVD method (2nd row, $v_{\text{cut}} = 0.01$), radial Wiener filter method (3rd row, $\alpha = 0.05$) and the transverse Wiener filter method (*bottom*, $\alpha = 0.05$) as described in the text.

Fig. 7. Reconstructions of a cluster at redshift $z = 0.25$ using noisy data with $n_g = 30$ galaxies per square arcminute. Shown above are the 2D projection of the reconstruction (*top left*), 1D line of sight plots for the four central lines of sight (*top right*) and smoothed 3D rendering of the reconstruction, as before.

an excellent correlation between the input density contrast and our reconstruction is seen. One line of sight exhibits a prominent high-redshift false detection; but this does not appear in the remaining three lines of sight, and the overall amplitude of the reconstruction is around 75% of the true value. The 3D rendering demonstrates that the noise in our reconstruction shows very little coherent structure (i.e. tends to be restricted to isolated pixels), and is largely low-level. Moreover, the cluster is incredibly well-localised in redshift space, with the smearing seen in the figure primarily arising from the applied smoothing.

Reconstruction with noisier data

The simulations described in Sect. 4.1 represent a fairly optimistic set of data. Realistically, one might expect a much lower number density of galaxies, so it is important to consider how our method fares with noisier data. To this end, weak lensing data were simulated for the same cluster as described above, but with $n_g = 30$ galaxies per square arcminute. This represents a factor of ~ 1.8 reduction in the signal to noise of our data. In order to account for this reduction in signal to noise, we increased our soft threshold parameter λ by a similar factor, taking $\lambda = 15$, whilst we again used $\epsilon = 3$.

The projections and 3D reconstruction obtained using these noisier data is shown in Fig. 7. The 2D projection shows a noisier reconstruction, with more false detections, and the central cluster appears to be less extended. This is to be expected, given the lower signal-to-noise ratio in the data. The line-of-sight plots again show more noise, with the peaks of some lines of sight being shifted slightly from the true redshift. However, they are well localised around the true peak, and their amplitudes match the true underlying density quite well.

The 3D rendering again shows a well-localised peak at the location of the cluster, albeit slightly extended along the line of sight. Note also that the false detections continue to appear as single pixel-scale detections, rather than as extended objects. As before, we believe that a fully 3D implementation of our algorithm will reduce the number of such false detections by searching for coherent structures of larger angular extent than a single pixel.

While improvements can be made to the noise model used, clearly, effective noise control can be obtained with noisy data with appropriate choice of parameters, and a very clean reconstruction obtained using our method. Here again, our method is seen to improve on the bias, smearing, and amplitude damping problems seen in the linear methods presented above, and the results at this noise level are of a comparable or better standard than the results from linear methods at higher signal-to-noise ratio, as presented above.

5.2. Improving the line-of-sight resolution

Given the success of our method at reconstructing lines of sight at the same resolution as the input data, it is interesting to consider whether we are able to improve on the output resolution of our reconstructions. It is also worthwhile to test our ability to detect clusters at higher redshifts than those considered above, given that compressed sensing is specifically designed to tackle underdetermined inverse problems. Indeed, noise-free simulations suggest that a resolution improvement of up to a factor of 4 in the redshift direction may be possible with this method.

Therefore, we generated clusters as before, with our data binned into $N_{sp} = 20$ redshift bins, but aimed to reconstruct onto $N_p = 25$ redshift bins. We furthermore consider clusters at redshifts of $z_{cl} = 0.2, 0.6$, and 1.0 . Given the changed reconstruction parameters, we modified our noise control parameters slightly and, for all results that follow, used $\lambda = 7.5$, $\epsilon = 2.8$.

Figure 8 shows the two-dimensional projection of the reconstruction and the 1D reconstruction of the four central lines of sight as before. Figure 9 shows the reconstructions of these haloes using our method as a 3D rendering. The 3D rendering is, as before, thresholded at $\delta_{rec} = 3$ and smoothed with a Gaussian of width $\sigma = 0.7$ pix.

Several features are immediately apparent. Firstly, all three clusters are clearly identified by our method. This is particularly impressive for the $z_{cl} = 1.0$ cluster, because linear methods have, thus far, been unable to reconstruct clusters at such a high redshift (see, e.g. Simon et al. 2009; VanderPlas et al. 2011). We note that this detection is dependent on the redshift distributions of sources, however, and the lack of detection in Simon et al. (2009) and VanderPlas et al. (2011) may be due, in part, to their choice of probability distribution. However, we also note that the background source density in our sample is highly diminished behind the $z_{cl} = 1.0$ cluster (32.5 arcmin^{-2}).

Again, the 3D renderings indicate that there is very little smearing of the reconstruction along the line of sight, in contrast with linear methods. This is also shown by the line-of-sight plots, in which the unsmoothed reconstructions show very localised structure. Furthermore, the reconstructions exhibit minimal redshift bias, and some lines of sight are seen to recover the amplitude of the density contrast without any notable damping.

However, again we see several prominent “hot pixels” or false detections along noisy lines of sight. These detections are more evident as the cluster moves to high redshift, and may be significantly larger than the expected density contrast of the

Fig. 8. Reconstructions of single clusters located at a redshift of $z_{cl} = 0.2$ (top row), $z_{cl} = 0.6$ (middle row) and $z_{cl} = 1.0$ (bottom row). As before, the left column shows the two-dimensional integrated projection of the reconstruction, while the right panel shows the input density contrast along the line of sight (solid line) and the 1D reconstruction along each of the four central lines of sight (dashed lines).

cluster. This is expected, because the number density of sources behind the lens diminishes.

We note that these false detections often manifest themselves at high redshift, arising out of the overfitting effect seen in Fig. 4. We also note that the false detections are very well localised in both angular and redshift space, and do not form coherent large structures, which makes them easily identifiable as false detections; they tend to be localised to isolated pixels. Because of this lack of coherence, we expect that a fully 3D implementation would suppress many of these false detections by aiming to detect coherent structure in three dimensions, and thereby seeking structures of larger extent than a single pixel.

5.3. More complex line-of-sight structure

Given the ability of our method to localise structure in redshift space, it is interesting to consider whether the algorithm is able to separate the lensing signal from two clusters located along the same line of sight. We considered three different cluster pairings – $z_{cl} = [0.2, 0.6]$, $z_{cl} = [0.2, 1.0]$, and $z_{cl} = [0.6, 1.0]$ – and the

results obtained by our method are shown in 3D rendering in Figure 10. We find that the reconstructions of individual lines of sight in this case is significantly more noisy than before.

The figure shows the $z_{cl} = [0.6, 1.0]$ cluster pairing to be reconstructed as a single, coherent structure smeared out between the two redshifts. In the other two cases, two distinct clusters are observed, but their redshifts are slightly biased, and a moderate amount of smearing in the redshift direction is seen. This smearing arises because individual lines of sight detect the structures at different redshifts; the aggregate effect is an elongation along the line of sight.

While these results are by no means as clean as the results obtained for single clusters, it is promising to note that we are able to detect the presence of more complex line of sight structure, even though the reconstruction of individual lines of sight are fairly noisy. Clearly, here, a fully 3D treatment of the data that takes into account the correlations between neighbouring lines of sight and that seeks to reconstruct coherent structures would offer improvements in this area.

6. Summary and conclusions

Current approaches to 3D weak lensing involve linear inversion, where a pseudo-inverse operator is constructed that incorporates prior constraints on the statistical distribution of the measurement noise and the underlying density. These methods are straightforward to construct and implement, make use of common tools, and are usually fairly fast. This makes them a convenient choice when approaching the 3D weak lensing problem.

However, reconstructions obtained in this way suffer from line-of-sight smearing, bias in the detected redshift of structures, and a damping of the reconstruction amplitude relative to the input. It has furthermore been noted (VanderPlas et al. 2011) that the reconstructions obtained using these techniques may be fundamentally limited regarding the resolution attainable along the line of sight, because of smearing effects resulting from these linear methods. In addition, these methods are unable to treat an underdetermined inversion, and therefore are limited in their output resolution by the resolution of the input data which, in turn, is limited by the measurement noise.

We have presented a new approach to 3D weak lensing reconstructions by considering the weak lensing problem to be an instance of compressed sensing, where the underlying structure we aim to reconstruct is sparsely represented in an appropriate dictionary. Under such a framework, we were able to consider underdetermined transformations, thereby relaxing the constraints on the resolution of the reconstruction obtained using our method.

We have reduced the problem to that of one-dimensional reconstructions along the line of sight. Whilst this is clearly not optimal because it throws away a lot of information, it allowed us to simplify the problem and to employ a particularly simple basis through which we imposed sparsity. We employed techniques recently developed in the area of convex optimisation to construct a robust reconstruction algorithm, and demonstrated that our method closely reproduces the position, radial extent and amplitude of simulated structures, with very little bias or smearing. This is a significant improvement over current linear methods. In addition, we have shown that our method produces clean results that demonstrate an improvement over linear methods, even when the signal-to-noise ratio in our data is reduced by a factor of ~ 2 in line with that expected from current lensing surveys.

Fig. 9. Reconstructions of single clusters located at a redshift of $z_{cl} = 0.2$ (*top*), $z_{cl} = 0.6$ (*middle*) and $z_{cl} = 1.0$ (*bottom*). The reconstructions are thresholded at $\delta = 3$ and smoothed with a Gaussian of width $\sigma = 0.7$ pix.

Fig. 10. Reconstructions of two clusters along the line of sight, located at a redshift of $z_{cl} = [0.2, 0.6]$ (*top*), $z_{cl} = [0.2, 1.0]$ (*middle*) and $z_{cl} = [0.6, 1.0]$ (*bottom*). The reconstructions are thresholded at $\delta = 3$ and smoothed with a Gaussian of width $\sigma = 0.7$ pix.

Furthermore, we demonstrated an ability to reconstruct clusters at higher redshifts than has been attainable using linear methods. Although our reconstructions exhibit false detections resulting from the noise, these noisy peaks do not form coherent structures, and are therefore well-localised and easily identifiable as noise peaks.

We have also tested the ability of our method to reconstruct multiple clusters along the line of sight. Whilst the reconstructions are noisy and exhibit a stronger redshift bias, damping, and smearing than that seen in the single cluster case, our method is seen to be sensitive enough to detect the presence of more than one structure along the line of sight. We expect that by improving our method, we will be able to more accurately reproduce these structures.

Clearly, this method holds a lot of promise; even in the one-dimensional implementation presented here, our results clearly show improvements to the bias, normalisation, and smearing problems seen with linear methods. There is much room for improvement, however, and a fully 3D implementation of the algorithm described here (which, as written, is entirely general) will be presented in future work. Such a treatment is expected to reduce the incidences of false detection in our reconstructions, the key to this improvement being the choice of an appropriate 3D dictionary with which to sparsify the solution.

In addition, we note that the simulations presented here are idealised as compared to real data. Most notably, we did not include any photometric redshift errors, which will be present in any real dataset. This is an important source of error in lensing measurements, and if the method presented above is to have any real value in reconstructing the matter distribution in a lensing survey, it must include treatment of these errors. This will be presented in a future work, included in the 3D implementation of the algorithm presented here.

The quality of data available for weak lensing measurements continues to improve, and the methods by which we measure the weak lensing shear are becoming ever more sophisticated. So, too, must the methods we use to analyse the data to reconstruct the underlying density. While linear methods appear to be limited in resolution, and offer biased estimators, it is clear that a nonlinear approach such as ours does not, and – in a fully 3D implementation – may therefore allow us to map the cosmic web in far greater detail than has previously been achieved.

Acknowledgements. We would like to thank Jalal Fadili for useful discussions, and Jake VanderPlas for supplying the code for the linear reconstructions and for helpful comments on an earlier version of this paper. We would also like to thank the anonymous referee for his/her useful comments and suggestions. This work is supported by the European Research Council grant SparseAstro (ERC-228261).

Appendix A: Solving the optimisation problem

In this section, we explain the development of the reconstruction algorithm. First, we will present an overview of some key concepts and results from convex optimisation, before introducing the primal-dual scheme chosen to solve Eq. (17), and finally discussing the convergence criterion for the algorithm. For a complete introduction to convex analysis, the reader is referred to Rockafellar (1970), Lemaréchal & Hiriart-Urruty (1996), Boyd & Vandenberghe (2004).

Notation and terminology

Let $\mathcal{H}$ be a real Hilbert space; in our case, a finite dimensional vector subspace of $\mathbb{R}^n$. We denote by $\|\cdot\|$ the norm associated

with the inner product in $\mathcal{H}$, I is the identity operator on $\mathcal{H}$, and $\|\cdot\|_p$ ($p \geq 1$) is the ℓ_p norm. A real-valued function f is *coercive* if $\lim_{\|x\| \rightarrow +\infty} f(x) = +\infty$, and is *proper* if its domain is non-empty:

$$\text{dom } f = \{x \in \mathcal{H} \mid f(x) < +\infty\} \neq \emptyset.$$

Lastly, $\Gamma_0(\mathcal{H})$ represents the class of all proper lower semicontinuous (lsc) convex functions from $\mathcal{H}$ to $(-\infty, +\infty]$.

A.1. Convex optimisation

The theory of convex optimisation aims to solve problems of the form

$$\hat{x} \in \underset{x \in \mathcal{H}}{\text{argmin}} F(x), \quad (\text{A.1})$$

where $F : \mathcal{H} \rightarrow \mathbb{R}$ is a convex function. If F represents a potential function, for example, then (A.1) will seek for the state of equilibrium, i.e. the state that minimises the potential.

Many algorithms exist for solving these problems, and an excellent introduction can be found in Boyd & Vandenberghe (2004). Recently, methods have been developed that exploit the decomposition of f into a sum of smaller convex functions. This can be very helpful if these functions show properties that are not preserved when the functions are summed.

For example, if we assume that F can be decomposed into a sum of K functions all in $\Gamma_0(\mathcal{H})$, then solving (A.1) is equivalent to solving

$$\underset{x \in \mathcal{H}}{\text{argmin}} \sum_{k=1}^K F_k(x). \quad (\text{A.2})$$

This formulation is interesting when the F_k functions show properties that are lost while considering the sum F , directly. In our case, we seek a formulation such that the functions we seek to minimise have proximal operators that are *simple*, and have closed form. While the sum F may not have this property, the functions F_k can be chosen such that they satisfy this condition.

A.2. Proximity operator

We may now begin to describe the proximal splitting algorithm used in this work. At the heart of the splitting framework is the notion of the *proximity operator*:

Definition 1 (Moreau 1962). Let $F \in \Gamma_0(\mathcal{H})$. Then, for every $x \in \mathcal{H}$, the function $y \mapsto F(y) + \|x - y\|_2^2 / 2$ achieves its infimum at a unique point denoted by $\text{prox}_F x$. The operator $\text{prox}_F : \mathcal{H} \rightarrow \mathcal{H}$ thus defined is the *proximity operator* of F .

Therefore, the proximity operator of the indicator function of a convex set C

$$I_C : x \mapsto \begin{cases} 0 & \text{if } x \in C, \\ \infty & \text{otherwise,} \end{cases}$$

is merely its orthogonal projector. Some key properties are presented in the following lemma:

Lemma 2 (Combettes & Wajs 2005).

- *Separability:* Let $F_k \in \Gamma_0(\mathcal{H})$, $k \in \{1, \dots, K\}$ and let $G : (x_k)_{1 \leq k \leq K} \mapsto \sum_k F_k(x_k)$. Then $\text{prox}_G = (\text{prox}_{F_k})_{1 \leq k \leq K}$.
- *Translation:* Let $F \in \Gamma_0(\mathcal{H})$ and $G \in \Gamma_0(\mathcal{H})$ such that $\forall x \in \mathcal{H}, G(x) = F(x - y)$, $y \in \mathcal{H}$. Then $\forall x \in \mathcal{H}$, $\text{prox}_G x = y + \text{prox}_F(x - y)$.

A.3. Primal-dual scheme

Consider the optimisation problem:

$$\hat{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{H}} G \circ \mathbf{U}(\mathbf{x}) + B(\mathbf{x}), \quad (\text{A.3})$$

where G and B are two convex, proper and lower semi-continuous functions and $\mathbf{U}$ is a linear bounded operator. Equation (A.3) can be seen as a special case of Eq. (A.2) where $K = 2$ and one of the functions contains a linear bounded operator. Then the algorithm 1 proposed in Chambolle & Pock (2011) will converge to the solution of (A.3), assuming that the proximity operators of G and B are easy to compute or known in closed form.

Algorithm 1: Primal-dual scheme for solving.

Parameters: The number of iterations N_{iter} , proximal steps $\sigma > 0$ and $\tau > 0$.

Initialisation:

$$\mathbf{x}^0 = \bar{\mathbf{x}}^0 = \mathbf{0}, \quad \xi^0 = \mathbf{0}.$$

Main iteration:

For $t = 0$ **to** $N_{\text{iter}} - 1$,

- **Dual step:** $\xi^{t+1} = (I - \sigma \operatorname{prox}_{G/\sigma})(\xi^t / \sigma + \mathbf{U}\bar{\mathbf{x}}^t)$.
- **Primal step:** $\mathbf{x}^{t+1} = \operatorname{prox}_{\tau B}(\mathbf{x}^t - \tau \mathbf{U}^* \xi^{t+1})$.
- **Update the coefficients estimate:** $\bar{\mathbf{x}}^{t+1} = 2\mathbf{x}^{t+1} - \mathbf{x}^t$.

End main iteration

Output: Final solution $\mathbf{x}^* = \mathbf{x}^{N_{\text{iter}}}$.

Adapting the arguments of (Chambolle & Pock 2011), convergence of the sequence $(\mathbf{x}^t)_{t \in \mathbb{N}}$ generated by algorithm 1 is ensured.

Proposition 3. Suppose that G and B are two convex, proper and lower semi-continuous functions. Let $\zeta = \|\mathbf{U}\|^2$, choose $\tau > 0$ and σ such that $\sigma\tau\zeta < 1$, and let $(\mathbf{x}^t)_{t \in \mathbb{N}}$ be that defined by Algorithm 1. Then, $(\mathbf{x}^t)_{t \in \mathbb{N}}$ converges to a (non-strict) global minimiser of Eq. (A.3).

A.4. Application to solution of Eq. (17)

First, we must rewrite the problem in Eq. (17) in an unconstrained form by replacing the constraints by the indicator functions of the corresponding constraint sets:

$$\min_{\mathbf{s} \in \mathbb{R}^n} \|\Phi^* \mathbf{s}\|_1 + \iota_{\ell_2(\varepsilon)}\left(\Sigma^{-\frac{1}{2}} \mathbf{d} - \Sigma^{-\frac{1}{2}} \mathbf{R} \mathbf{s}\right) + \iota_C(\mathbf{s}), \quad (\text{A.4})$$

where $\iota_{\ell_2(\varepsilon)}$ is the indicatrice function of the ℓ_2 -ball of radius ε and ι_C the indicatrice function of a closed convex set C .

Notice that (A.4) can be expressed in the form of (A.3) with

$$G(\mathbf{x}, \mathbf{y}) = \|\mathbf{x}\|_1 + \iota_{\ell_2(\varepsilon)}(\mathbf{y} - \Sigma^{-\frac{1}{2}} \mathbf{d}), \quad (\text{A.5})$$

$$B(\mathbf{x}) = \iota_C(\mathbf{x}), \quad (\text{A.6})$$

$$\mathbf{U} = \begin{pmatrix} \Phi^* \\ \Sigma^{-\frac{1}{2}} \mathbf{Q} \end{pmatrix}. \quad (\text{A.7})$$

We now need only to apply algorithm 1 to compute the solution. This requires computation of the three proximity operators, which are given by the following proposition:

Proposition 4. Let $F \in \Gamma_0(\mathcal{H})$. Then,

- if $F : \mathbf{x} \mapsto \|\mathbf{x}\|_1$, then its associated proximity operator, $\operatorname{prox}_{\lambda F}$, is the component-wise soft-thresholding operator with threshold λ as defined by Eq. (B.1);

- if $F : \mathbf{x} \mapsto \iota_C(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in C, \\ \infty & \text{else,} \end{cases}$ then its associated proximity operator, prox_F , is the Euclidian projector to the convex set C , Prj_C ;
 - if $F : \mathbf{x} \mapsto \iota_{\ell_2(\varepsilon)}(\mathbf{x}) = \begin{cases} 0 & \text{if } \|\mathbf{x}\|_2 \leq \varepsilon, \\ \infty & \text{else,} \end{cases}$ then its associated proximity operator, prox_F , is the projector onto the ℓ_2 -ball with radius ε defined as
- $$\operatorname{prox}_F \mathbf{x} = \begin{cases} \mathbf{x} & \text{if } \|\mathbf{x}\|_2 \leq \varepsilon, \\ \frac{\varepsilon}{\|\mathbf{x}\|_2} \mathbf{x} & \text{else.} \end{cases}$$

The application of proposition 4 and lemma 2 to the optimisation problem in Eq. (A.4), using the primal-dual scheme in algorithm 1, yields the reconstruction algorithm expressed in algorithm 2.

Appendix B: Practical considerations

Application of the method described in Appendix A to the specific case of 3D lensing leads us to algorithm 2, where one can note the projection over the two constraints described above (data fidelity and minimum value of the solution), and the operator $\mathcal{S}t_\lambda$, which imposes a soft threshold at a level of λ/ω as

$$\mathcal{S}t_\lambda : \alpha \mapsto (g_\lambda(\alpha_i))_{1 \leq i \leq L}, \quad (\text{B.1})$$

$$g_\lambda : \eta \mapsto \begin{cases} (|\eta| - \lambda) \operatorname{sign}(\eta) & |\eta| > \lambda \\ 0 & \text{otherwise.} \end{cases}$$

B.1. Convergence criteria

The main difficulty with the primal-dual scheme described above – indeed, with any iterative algorithm – is to define an appropriate convergence criterion. In this case, the difference between two successive iterates $\mathbf{x}^t$ and $\mathbf{x}^{t+1}$ is not bounded. However, the partial primal-dual gap $\mathbb{G}_{\text{pd}}$, defined by

$$\mathbb{G}_{\text{pd}}(\mathbf{x}, \mathbf{y}) = \max_{\mathbf{y}'} \langle \mathbf{y}', \mathbf{U} \mathbf{x} \rangle - G^*(\mathbf{y}') + B(\mathbf{x}) - \min_{\mathbf{x}'} \langle \mathbf{y}, \mathbf{U} \mathbf{x}' \rangle - G^*(\mathbf{y}) + B(\mathbf{x}'), \quad (\text{B.2})$$

(when solving (A.3)) is bounded. Here, G^* is the convex conjugate of convex function G :

$$G^* : \mathbf{x} \mapsto \max_{\mathbf{y}} \langle \mathbf{x}, \mathbf{y} \rangle - G(\mathbf{y}). \quad (\text{B.3})$$

Let us define two variables from the sequences $(\mathbf{x}^t)_t$ and $(\xi^t)_t$ produced by algorithm 1:

$$\mathbf{x}^N = \frac{1}{N} \sum_{t=0}^{N-1} \mathbf{x}^t \quad \text{and} \quad \xi^N = \frac{1}{N} \sum_{t=0}^{N-1} \xi^t,$$

which are the accumulation variables at iteration $N - 1$. Chambolle & Pock (2011) have shown that the sequence defined by $(\mathbb{G}_{\text{pd}}(\mathbf{x}^N, \xi^N))_{N \in \mathbb{N}}$ is bounded, and decreases at a rate of $\mathcal{O}(1/t)$, where t is the iteration number.

In order to use (B.2) in our context, the two indicatrice functions inside (A.4) are not considered, because they play little role. Therefore, in our case, we may rewrite $\mathbb{G}_{\text{pd}}(\mathbf{x}, \mathbf{y})$ as:

$$\mathbb{G}_{\text{pd}}(\mathbf{x}, \mathbf{y}) \approx \max_{\mathbf{y}'} \langle \mathbf{y}', \mathbf{U} \mathbf{x} \rangle - G^*(\mathbf{y}') - G^*(\mathbf{y}), \quad (\text{B.4})$$

$$\approx G(\mathbf{U} \mathbf{x}) - G^*(\mathbf{y}), \quad (\text{B.5})$$

$$\approx \lambda \|\Phi^* \mathbf{x}\|_1. \quad (\text{B.6})$$

Algorithm 2: Nonlinear iterative algorithm for solving Eq. (17).

Parameters: Choose $\omega, \tau > 0$ such that $\omega\tau\Theta^2 \leq 1$ where

$\Theta \equiv \left\| \Sigma^{-\frac{1}{2}} \mathbf{R} \right\|_2 + \|\Phi^*\|_2$, where $\|\mathbf{R}\|_2$ is the spectral norm of the operator $\mathbf{R}$.

Initialization:

$\mathbf{y}_1^0 = \mathbf{y}_2^0 = \mathbf{0}$, $\hat{\mathbf{s}}^0 = \mathbf{0}$, $\mathbf{x}^0 = \mathbf{0}$.

N = number of elements in $\mathbf{d}$.

$\hat{\mathbf{d}} = \Sigma^{-\frac{1}{2}} \mathbf{d}$.

Choose $\lambda > 0$.

Main iteration:

For $n = 0$ **to** $N_{\text{iter}} - 2$,

– Initialise auxiliary variables:

$$\mathbf{t}_1^n = \mathbf{y}_1^n + \omega \Phi^* \hat{\mathbf{s}}^n.$$

$$\mathbf{t}_2^n = \mathbf{y}_2^n + \omega \Sigma^{-\frac{1}{2}} \mathbf{R} \hat{\mathbf{s}}^n.$$

– Sparsity-promoting penalty:

$$\mathbf{y}_1^{n+1} = \mathbf{t}_1^n - \omega S_{t/\omega}(\mathbf{t}_1^n/\omega).$$

– Data fidelity term:

$$\mathbf{y}_2^{n+1} = \mathbf{t}_2^n - \omega(\hat{\mathbf{d}} + \zeta)$$

$$\left\{ \begin{array}{l} \mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega \quad \text{if } \|\mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega\|_2 < \epsilon\sqrt{N}, \\ \epsilon\sqrt{N}(\mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega) / \|\mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega\|_2 \quad \text{otherwise.} \end{array} \right.$$

$$\zeta = \left\{ \begin{array}{l} \epsilon\sqrt{N}(\mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega) / \|\mathbf{t}_2^n/\omega - \hat{\mathbf{d}}/\omega\|_2 \quad \text{otherwise.} \end{array} \right.$$

– Projection on the convex set C :

$$\mathbf{x}^{n+1} = \text{Prj}_C \left[\mathbf{x}^n - \tau \left(\Phi \mathbf{y}_1^{n+1} + \mathbf{R}^\dagger \Sigma^{\frac{1}{2}} \mathbf{y}_2^{n+1} \right) \right].$$

– Update current estimate:

$$\hat{\mathbf{s}}^{n+1} = 2\mathbf{x}^{n+1} - \mathbf{x}^n.$$

End main iteration

We determine the algorithm to have converged when

$$\Delta \mathfrak{G}_{\text{pd}}^N \equiv \frac{\mathfrak{G}_{\text{pd}}(\mathbf{x}^N, \xi^N) - \mathfrak{G}_{\text{pd}}(\mathbf{x}^{N-1}, \xi^{N-1})}{\mathfrak{G}_{\text{pd}}(\mathbf{x}^{N-1}, \xi^{N-1})} < \alpha, \quad (\text{B.7})$$

where the appropriate α can be determined from simulations, and is dependent on a tradeoff between the desired level of accuracy of the reconstructed data and the time taken to complete the reconstruction. Choosing α to be large will result in an estimate of the solution that may be some distance away from the solution that would be obtained if absolute convergence were reached, but which is obtained in a small number of iterations.

Figure B.1 shows a characteristic example of the function $\Delta \mathfrak{G}_{\text{pd}}^N$ as a function of iteration number N for the simulations described in Sect. 5. This function is clearly a smooth, largely steadily decreasing function of iteration number, and thus an appropriate choice for defining convergence. Note that the curve shows some oscillations with iteration number. These oscillations arise on lines of sight where the noise results in the algorithm having difficulty fitting the data within the constraints. With appropriate choice of parameters, these oscillations are relatively small and the curve eventually becomes smooth.

We find from experimentation that $\alpha = 10^{-6}$ yields a solution that is sufficiently accurate for our purposes, and is sufficient to largely remove the oscillations, therefore we set this to be our

Fig. B.1. The function $\Delta \mathfrak{G}_{\text{pd}}^N$, as defined in the text, as a function of iteration number.

threshold for all reconstructions presented above. In addition, we require that each line of sight undergoes at least 1500 iterations, to avoid misidentification of convergence due to early oscillations of $\Delta \mathfrak{G}_{\text{pd}}^N$, specifically a strong dip seen in this curve at the start of iteration.

It is possible, along certain lines of sight, for the estimate of the solution to remain constant at zero owing to the soft thresholding while $\Delta \mathfrak{G}_{\text{pd}}^N$ still varies. To account for this, if the current cumulative estimate $\mathbf{x}^N$ does not vary for 200 iterations, we assume that convergence has been reached for that line of sight.

We now consider several practical points involved in the implementation of this algorithm.

B.2. Choice of step sizes ω , τ

ω and τ control the step size in the evolution of the algorithm, and are required to be positive and to satisfy the inequality

$$\omega\tau\Theta^2 \leq 1 \quad (\text{B.8})$$

where $\Theta \equiv \left\| \Sigma^{-\frac{1}{2}} \mathbf{Q} \right\|_2 + \|\Phi^*\|_2$ is the sum of the ℓ_2 norms of the operators used in the algorithm. This sum is dominated by the second term, because the elements of the $\mathbf{Q}$ lensing efficiency matrix are small. We have chosen a δ -function dictionary, therefore the application of the transformation Φ^* represents a multiplication by the identity matrix.

Accordingly, we have $\Theta \sim \|\Phi^*\|_2 \sim 1$, which implies that $\omega\tau \sim 1$ is appropriate. For all results that follow, we choose $\omega = \tau = 1$. Lower values of these parameters may be used with little effect on the resulting solution. However, the lower the chosen values, the smaller the steps taken in each iteration of the algorithm. This results in a slower convergence than seen with higher values of ω and τ . Choosing higher values than implied by Eq. (B.8) is not advised, because convergence of the algorithm is not guaranteed in this case; the algorithm may give rise to a strongly oscillating estimator, and the resulting solution may not be the one that fits the data best given the prior constraints.

References

- Aitken, R. G. 1934, Proc. R. Soc. Edinb., 55, 42
- Albrecht, A., Bernstein, G., Cahn, R., et al. 2006 [arXiv:astro-ph/0609591]
- Bacon, D. J., & Taylor, A. N. 2003, MNRAS, 344, 1307
- Barbey, N., Sauvage, M., Starck, J.-L., Ottensamer, R., & Chaniai, P. 2011, A&A, 527, A102

- Bobin, J., Starck, J.-L., & Ottensamer, R. 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 718
- Boyd, S., & Vandenberghe, L. 2004, *Convex Optimization* (Cambridge University Press)
- Candes, E. J., & Plan, Y. 2010, *A Probabilistic and RIPless Theory of Compressed Sensing*, submitted
- Candès, E., & Tao, T. 2006, *IEEE Transactions on Information Theory*, 52, 5406
- Castro, P. G., Heavens, A. F., & Kitching, T. D. 2005, *Phys. Rev. D*, 72, 023516
- Chambolle, A., & Pock, T. 2011, *J. Mathem. Imag. Vision*, 40, 120
- Combettes, P. L., & Wajs, V. R. 2005, *SIAM Multiscale Model. Simul.*, 4, 1168
- Donoho, D. 2006, *IEEE Transactions on Information Theory*, 52, 1289
- Fadili, M., & Starck, J.-L. 2009, in *Proc. International Conference on Image Processing, ICIP 2009*, 7–10 November, Cairo, Egypt (IEEE), 1461
- Hoekstra, H., & Jain, B. 2008, *Ann. Rev. Nucl. Par. Sci.*, 58, 99
- Hu, W., & Keeton, C. R. 2002, *Phys. Rev. D*, 66, 063506
- Ivezic, Z., Tyson, J. A., Allsman, R., et al. 2008 [[arXiv:0805.2366](#)]
- Kitching, T. D., Heavens, A. F., & Miller, L. 2011, *MNRAS*, 426
- Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, *ApJS*, 192, 16
- Lemaréchal, C., & Hiriart-Urruty, J.-B. 1996, *Convex Analysis and Minimization Algorithms I*, 2nd ed. (Springer)
- Levy, D., & Brustein, R. 2009, *JCAP*, 6, 26
- Li, F., Cornwell, T. J., & de Hoog, F. 2011, *A&A*, 528, A31
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009 [[arXiv:0912.0201](#)]
- Ma, Z., Hu, W., & Huterer, D. 2006, *ApJ*, 636, 21
- Massey, R., Rhodes, J., Ellis, R., et al. 2007a, *Nature*, 445, 286
- Massey, R., Rhodes, J., Leauthaud, A., et al. 2007b, *ApJS*, 172, 239
- Moreau, J.-J. 1962, *CRAS Sér. A Math.*, 255, 2897
- Munshi, D., Valageas, P., van Waerbeke, L., & Heavens, A. 2008, *Phys. Rep.*, 462, 67
- Peacock, J. A., Schneider, P., Efstathiou, G., et al. 2006, *ESA-ESO Working Group on Fundamental Cosmology*, Tech. rep., Ed. University
- Refregier, A., Amara, A., Kitching, T. D., et al. 2010 [[arXiv:1001.0061](#)]
- Rockafellar, R. 1970, *Convex analysis* (Princeton University Press)
- Schneider, P. 2006, *Weak Gravitational Lensing* (Springer), 269
- Simon, P., Taylor, A. N., & Hartlap, J. 2009, *MNRAS*, 399, 48
- Simon, P., Heymans, C., Schrabbach, T., et al. 2011, *MNRAS*, 1789
- Starck, J.-L., Murtagh, F., & Fadili, J. M. 2010, *Sparse Image and Signal Processing* (Cambridge University Press)
- Taylor, A. N., Bacon, D. J., Gray, M. E., et al. 2004, *MNRAS*, 353, 1176
- Taylor, A. N., Kitching, T. D., Bacon, D. J., & Heavens, A. F. 2007, *MNRAS*, 374, 1377
- Tegmark, M. 1997, *ApJ*, 480, L87
- Tegmark, M., de Oliveira-Costa, A., Devlin, M. J., et al. 1997, *ApJ*, 474, L77
- Van Waerbeke, L., & Mellier, Y. 2003 [[arXiv:astro-ph/0305089](#)]
- VanderPlas, J. T., Connolly, A. J., Jain, B., & Jarvis, M. 2011, *ApJ*, 727, 118
- Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., & Vanderghenst, P. 2009, *MNRAS*, 395, 1733