

SPARSE BAYESIAN MASS-MAPPING USING TRANS-DIMENSIONAL MCMC

AUGUSTIN MARIGNIER,^{1,2†} THOMAS KITCHING,¹ JASON D. McEWEN¹ ANA M. G. FERREIRA^{2,3}

¹Mullard Space Science Laboratory, University College London, Dorking RH5 6NT, UK

²Department of Earth Sciences, University College London, London WC1E 6BT, UK

³CERIS, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

Version November 24, 2022

ABSTRACT

Uncertainty quantification is a crucial step of cosmological mass-mapping that is often ignored. Suggested methods are typically only approximate or make strong assumptions of Gaussianity of the shear field. Probabilistic sampling methods, such as Markov chain Monte Carlo (MCMC), draw samples from a probability distribution, allowing for full and flexible uncertainty quantification, however these methods are notoriously slow and struggle in the high-dimensional parameter spaces of imaging problems. In this work we use, for the first time, a trans-dimensional MCMC sampler for mass-mapping, promoting sparsity in a wavelet basis. This sampler gradually grows the parameter space as required by the data, exploiting the extremely sparse nature of mass maps in wavelet space. The wavelet coefficients are arranged in a tree-like structure, which adds finer scale detail as the parameter space grows. We demonstrate the trans-dimensional sampler on galaxy cluster-scale images where the planar modelling approximation is valid. In high-resolution experiments, this method produces naturally parsimonious solutions, requiring less than 1% of the potential maximum number of wavelet coefficients and still producing a good fit to the observed data. In the presence of noisy data, trans-dimensional MCMC produces a better reconstruction of mass-maps than the standard smoothed Kaiser-Squires method, with the addition that uncertainties are fully quantified. This opens up the possibility for new mass maps and inferences about the nature of dark matter using the new high-resolution data from upcoming weak lensing surveys such as *Euclid*.

1. INTRODUCTION

Gravitational lensing is the phenomenon where light from distant objects is distorted by the density field between the object and the observer, resulting in stunning images of warped, sheared, magnified and multiplied galaxies. Lensing is due to both light and dark forms of matter, making it a promising probe of the nature of dark matter (Heavens 2009). Weak lensing is the regime of small distortions, the effect of which on distant galaxies can be described as a magnification due to a convergence field κ and a perturbation of intrinsic ellipticity due to a shear field γ .

The convergence field κ is the integrated total mass density along the line of sight (Bartelmann & Schneider 2001; Dodelson 2017), and hence is a measure of the matter overdensity field. The density field contains both Gaussian and non-Gaussian structures. The Gaussian structures of the cosmological initial conditions evolve under the non-linear influence of gravity to produce non-Gaussian structures. As such, it is essential that methods to build maps of the convergence field (mass maps) contain Gaussian and non-Gaussian structures. With non-Gaussian structures, the higher-order statistics such as Minkowski functions and bispectrum (Munshi & Coles 2017) of the mass maps can be measured in addition to the usual two-point correlation function (Peebles 1980). From these statistics, one can compare with the predictions of the statistical density distributions of Λ CDM or other cosmologies. Further, one can compare the total density distribution in colliding cluster systems with the observed distribution of baryonic matter to infer the presence and collisional properties of dark matter (e.g. Clowe et al. 2006; Harvey et al. 2015). These sorts of comparisons would be much better informed if uncertainties were also quantified with the mass maps.

Unfortunately, κ cannot be directly observed, and must be constrained by an ill-posed inverse problem of observed ellipticities modelling the shear field γ . There have been many methods proposed to solve this inverse problem. The standard method is the Kaiser-Squires (KS, Kaiser & Squires 1993), which is a simple Fourier space linear operator. However, it is well known that this method is not robust to noise and missing data, and is often smoothed to mitigate the effects of noise at the expense of small-scale non-Gaussian structures. Despite this the mass maps recovered with KS are generally reliable, and the method has been extended to the sphere (Wallis et al. 2021) for use in wide-field surveys (Van Waerbeke et al. 2013; Chang et al. 2018; Jeffrey et al. 2018, 2021; Price et al. 2020b). Recently, there has been a push for methods with sparsity-promoting regularisation to preserve non-Gaussianities and deal with irregular data sampling (e.g. Lanusse et al. 2016; Price et al. 2021; Starck et al. 2021), where sparsity is typically promoted in a wavelet basis.

Despite all these methods, uncertainty quantification for mass-mapping remains a difficult task. Common approaches for uncertainty quantification involve computationally expensive posterior sampling by, for example, Markov chain Monte Carlo (MCMC) methods. In the weak lensing context, however, only recently have these methods been used for mass-mapping (e.g. Alsing et al. 2015; Price et al. 2020a; Porqueres et al. 2021; Fiedorowicz et al. 2022; Remy et al. 2022). This is explained by imaging inverse problems being very high-dimensional, resulting in difficulties in converging to the posterior. To make the sampling of the posterior computationally tractable, the slope of the posterior is often exploited to identify the high-density regions, for example in Hamiltonian Monte Carlo (Neal 2012) or proximal MCMC (Pereyra 2016). Both of these approaches have now been used for mass-mapping (Price et al. 2020a; Fiedorowicz et al. 2022; Remy et al. 2022;

† Correspondence address: augustin.marignier.14@ucl.ac.uk

Marignier et al. 2022). Various priors have been used, including Gaussian (Alsing et al. 2015), sparsity (Price et al. 2020a; Marignier et al. 2022), physically informed power spectrum (Porqueres et al. 2021), log-normal (Fiedorowicz et al. 2022) and learnt (Remy et al. 2022) priors.

The purpose of this work is to introduce a new efficient MCMC sampling method for mass-mapping. We use a trans-dimensional MCMC sampler to gradually grow the parameter space as required by the data. Trans-dimensional inversions, where the number of model parameters itself a parameter to be determined (Geyer & Møller 1994; Green 1995), are now extremely common in geophysical studies (e.g. Malinverno 2002; Gallagher 2012; Bodin & Sambridge 2009; Minsley 2011; Tkalčić et al. 2013; Piana Agostinetti et al. 2015; Hawkins & Sambridge 2015; Burdick et al. 2019), have been used sporadically in astronomy (e.g. Cornish & Littenberg 2007; Karnesis et al. 2014; Feder & Daylan 2018) and have yet to be used for mass-mapping. The particular sampler we use was proposed by Hawkins & Sambridge (2015) for geophysical imaging, and uses a wavelet tree parameterisation allowing for multi-scale analysis and sparsity-promotion. The appeal of this method is its ability to significantly reduce the effective parameter space. For example, Hawkins & Sambridge (2015) showed a 3D seismic tomography example with a parameter space of at most 524,288 parameters, and their method produced a solution with around 500 parameters. This sort of efficiency in sampling the parameter space will be essential for high-resolution mass-maps as the sky-fraction covered by upcoming surveys increases. Other methods for sampling such high parameter spaces typically require gradient or proximal mapping calculations (e.g. Neal 2012; Pereyra 2017), which are avoided with this trans-dimensional method.

In this work we use the trans-dimensional tree MCMC sampler of Hawkins & Sambridge (2015) for Bayesian mass-mapping with full uncertainty quantification. We exploit the multi-scale nature of the wavelet parameterisation to impose a scale-dependent sparsity-promoting prior using a generalised Gaussian distribution (GGD), as opposed to the more common Laplace distribution. The rationale for this is that to first order larger scale structures in mass maps are more Gaussian than the smaller scale details. A similar rationale was used by Starck et al. (2021) for mass mapping, although our approach is more similar to that used by McEwen et al. (2017) for the analysis of cosmic strings in the cosmic microwave background. We demonstrate this method on the recovery of cluster-scale convergence maps from N-body simulations.

This article is organised as follows. In Section 2 we outline the theoretical background of weak lensing and Bayesian inversion. In Section 3 we give details of the trans-dimensional tree inversion. Section 4 outlines the specificities of the mass-mapping inverse problem. Section 5 presents our results on simulation data, which we discuss further in Section 6. We conclude in Section 7.

2. BACKGROUND

2.1. Weak Lensing and Mass-Mapping

Gravitational lensing describes the deflection of photons from distant sources as they pass by regions of local gravitational potential variations caused by the local over or under density of matter acting as a lens. The weak lensing effect can be described in terms of a 3D lensing potential $\psi(\mathbf{r}) = \psi(\theta_1, \theta_2, r)$, which is the integral along the line of

sight of the Newtonian gravitational potential $\Phi(\mathbf{r})$

$$\psi(\mathbf{r}) = \frac{2}{c^2} \int_0^r \frac{f_K(r-r')}{f_K(r)f_K(r')} \psi(r'\theta_1, r'\theta_2, r) dr'. \quad (1)$$

Here, r is the comoving distance from the observer to the source, (θ_1, θ_2) are angular positions on the sky, c is the speed of light in a vacuum, and $f_K(r)$ is the comoving angular separation in a cosmology with curvature K , which we take to be 0 (i.e. a flat universe). The gravitational potential is related to the fractional matter over-density field $\delta(\mathbf{r})$ by Poisson's equation

$$\nabla^2 \Phi = \frac{3\Omega_m H_0^2}{2c^2 a(t)} \delta(\mathbf{r}) \quad (2)$$

where Ω_m is the current matter density parameter, H_0 is the current Hubble constant, and $a(t)$ is the scale factor. The aim of mass-mapping is to infer the convergence field κ from galaxy distortion measurements representing the shear field γ . The convergence represents the total density perturbation along the line of sight. The convergence and the shear are defined as gradients of the lensing field, and, as previously derived, are related in Fourier space by

$$\tilde{\gamma}(\mathbf{l}) = \frac{l_1^2 - l_2^2 + 2il_1 l_2}{l_1^2 + l_2^2} \tilde{\kappa}(\mathbf{l}) \quad (3)$$

where (l_1, l_2) is the conjugate variable of (θ_1, θ_2) . This defines the Fourier space KS kernel (Kaiser & Squires 1993). Note that it is undefined at the origin $l_1 = l_2 = 0$, which corresponds to the mass-sheet degeneracy. The mass-sheet degeneracy implies that κ can only be determined up to an additive constant. Physically, this is because a constant surface mass density does not cause any shear (Bartelmann & Schneider 2001).

2.2. Bayesian Inversion

From Bayes' Theorem, the *posterior* probability distribution describing the probability that the underlying convergence field producing observed complex shear measurements γ is described by some discretised field κ , is given by

$$p(\kappa|\gamma) \propto p(\gamma|\kappa)p(\kappa) \quad (4)$$

where $p(\gamma|\kappa)$ is the likelihood function describing data fidelity between the observed γ and predictions generated from κ , and $p(\kappa)$ is a distribution describing any prior belief about the convergence field, e.g. that it is sparse in a wavelet basis.

In standard MCMC, a sampler will start at some model of convergence κ_i and propose some new model κ'_i based on a symmetric proposal distribution $q(\kappa'_i|\kappa_i)$. This distribution is symmetric in that $q(\kappa'_i|\kappa_i) = q(\kappa_i|\kappa'_i)$, i.e. the probability of proposing κ'_i when starting at κ_i is the same as the probability of proposing κ_i when starting at κ'_i . The proposed model is either accepted ($\kappa_{i+1} = \kappa'_i$) or rejected ($\kappa_{i+1} = \kappa_i$) according to an acceptance criterion. A common choice of acceptance criterion is known as the Metropolis-Hastings (MH) criterion

$$\alpha(\kappa', \kappa) = \min \left\{ 1, \frac{p(\kappa')p(\gamma|\kappa')q(\kappa_i|\kappa'_i)}{p(\kappa_i)p(\gamma|\kappa_i)q(\kappa'_i|\kappa_i)} \right\} \quad (5)$$

Repeating this many thousands of times produces a chain of samples which, thanks to the acceptance criterion, is guaranteed to converge to the posterior distribution, giving the solution of the inverse problem. A summary convergence map (e.g. the sample mean) can be calculated, as can any measure

of uncertainty. This is a great strength of sampling methods, such as MCMC, although it comes at great computational cost.

2.3. Wavelet Transform

The wavelet transform is a popular tool in image processing, having been popularised in the late 20th century (Mallat 1989; Daubechies 1992). Since then, the theory of compressed sensing (Donoho 2006; Candès et al. 2011) showed that sparse signals can be accurately recovered from incomplete data, leading many studies in different fields using sparsity-promoting priors or regularisation in a wavelet space, including weak lensing studies (e.g. Lanusse et al. 2016; Price et al. 2021; Starck et al. 2021).

Wavelets are rapidly-decaying, wave-like functions that exist for a finite amount of time or space and form a basis in which typical signals tend to be sparse. There are many different families of wavelets, the details of which are beyond the scope of this work. The discrete wavelet transform can be seen as high and low-pass convolutions using scaled and translated versions of the “mother” wavelet, resulting in a separation of localised information at different length scales. Typically these high and low pass filters extract the high and low half, respectively, of the frequencies in the image. The outputs of the high-pass filters are the detailed wavelet coefficients, and the low-pass filters give the approximation coefficients. Repeating the process on the approximation coefficients returns detailed wavelet coefficients at a slightly larger scale than the previous pass. After each pass the wavelet coefficients are downsampled by a factor of two, as only half the frequency content from the previous scale remains. The final set of approximation coefficients are known as the coefficients of the scaling function. The scaling function is distinct from the wavelets as it accounts for the lowest frequency content, in particular at 0 frequency.

For the 2D discrete wavelet transform, convolutions are performed horizontally and vertically, giving three sets of detailed wavelet coefficients at each scale: horizontal, vertical and diagonal details. For an image with dimensions $2^J \times 2^J$, for some integer J , after J passes one is left with a single scaling coefficient and $2^J \times 2^J - 1$ wavelet coefficients. The scaling coefficient represents the overall background mean of the original image. In the following we denote the various wavelet scales by the variable j , with $j = 0$ representing the largest scale (low frequency) and $j = J$ representing the smallest scale (high frequency).

3. TRANS-DIMENSIONAL TREE INVERSION

In trans-dimensional MCMC, the parameterisation and the dimensionality of κ is allowed to change (Green 1995). For example, the number of wavelet coefficients describing κ is variable. This allows the available data to determine the complexity of the solution, rather than having this set *a priori*. Instinctively, it might seem that such a trans-dimensional inversion will naturally tend to more complex models, as one can obtain arbitrarily better fits to data by adding more and more model parameters. It has been shown however that results are generally parsimonious. To generalise the MH acceptance criterion, the determinant of a Jacobian matrix describing the transformation from one parameter space to another is included, i.e

$$\alpha(\kappa', \kappa) = \min \left\{ 1, \frac{p(\kappa'_i) p(\gamma' | \kappa'_i) q(\kappa_i | \kappa'_i)}{p(\kappa_i) p(\gamma | \kappa_i) q(\kappa'_i | \kappa_i)} | \mathcal{J} | \right\}. \quad (6)$$

Significantly though, as we outline in the following section, with simple choices about the parameterisation and proposal distribution the determinant of the Jacobian is 1, effectively keeping the original acceptance criterion.

The trans-dimensional sampler we use here is also known as a birth/death sampler, where model parameters are gradually added or removed, as well have their values perturbed as in standard MCMC (Green 1995). In this work, we use a parameterisation first proposed by Hawkins & Sambridge (2015) for geophysical inversions. Described as a tree structure, nodes can be added and removed from the tree via the birth and death proposals, respectively. The nodes are arranged such that birth proposals add more detail to the model. We give some details of the structure and the prior and proposal distributions here, although for brevity we only describe the choices made for our mass-mapping application. For more details, we refer the reader to Hawkins & Sambridge (2015). The aim of using this tree structure in the mass-mapping context is to gradually add more and more small-scale detail.

3.1. The wavelet tree model

We parameterise our convergence maps using a set of k wavelet coefficients \varkappa in a tree structure, denoted by \mathcal{T}_k . The root of the tree is the scaling function coefficient, a single-pixel representation of the image. The tree then branches down to smaller scale structure. This is shown schematically in Figure 1. The k coefficients of the current tree are indicated by black dots and their connection in the tree is indicated by the red arrows. The maximum depth of the tree is the J^{th} wavelet scale, chosen based on the desired output resolution. Our model space vector κ is then given by

$$\kappa = \langle \mathcal{T}_k, \{\varkappa_i : i \in 1, \dots, k\} \rangle, \quad (7)$$

where $\langle \cdot, \dots, \cdot \rangle$ denotes a vector and i is a unique index for all the coefficients in the tree.

3.2. The prior

The prior $p(\kappa)$ consists of three parts: (1) the probability of the number of coefficients, $p(k)$; (2) the probability that k coefficients are arranged in the tree \mathcal{T}_k , $p(\mathcal{T}_k | k)$; and (3) the prior on the wavelet coefficient values $p(\varkappa_i | \mathcal{T}_k, k)$.

For the probability of k we choose a simple uniform prior

$$p(k) = \frac{1}{k_{\max}} \quad (8)$$

where we take $k_{\max} = 2^J$, i.e. the theoretical number of wavelet coefficients with J wavelet scales. This prior is effectively uninformative, as we have no initial prior belief as to how many wavelet coefficients should be in the tree. This choice also simplifies the calculation of the proposal probabilities, which we describe in detail in subsection 3.3.

Similarly, we use a uniform prior for the arrangement of \mathcal{T}_k . Our tree is *restricted and heterogeneous*. Restricted trees have a maximum depth, in our case the number of wavelet scales J . Heterogeneous trees vary the number of child coefficients that each parent coefficient has. For the 2D wavelet tree, the scaling coefficient ($j = 0$) has 3 children, and then each coefficient from $j = 1$ onwards has 4 children (Figure 1). Our prior is then

$$p(\mathcal{T}_k | k) = \frac{1}{\mathcal{N}_{k,J}}, \quad (9)$$

where $\mathcal{N}_{k,J}$ is the number of valid heterogeneous tree configurations of k coefficients restricted to depth J . This is calculated

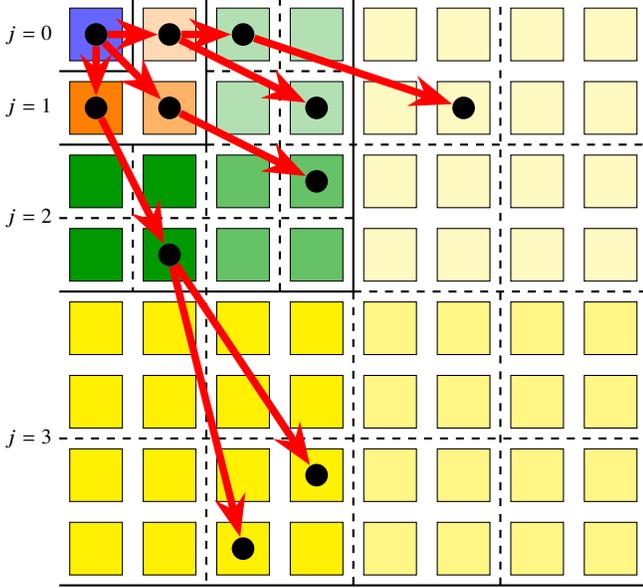


FIG. 1.— The 2D wavelet tree parameterisation for a tree of maximum depth 3. Each coloured square corresponds to a wavelet coefficient, although only those with black dots form part of the current tree (red arrows). This current tree has $k = 11$ coefficients. The coefficients of each wavelet scale ($j = \{0, 1, 2, 3\}$) are separated by solid black lines and grouped by colour (blues, oranges, greens, yellows). The blue coefficient at the root of the tree is the scaling coefficient. Beneath this I have 3 wavelet coefficients at scale $j = 1$. Beneath each of these three coefficients I have 4 coefficients at scale $j = 2$, where the relation is represented by the colour shade e.g. the 4 dark green coefficients in $j = 2$ are beneath the dark orange coefficient in $j = 1$. From each of the $j = 2$ coefficients I then get another 4 coefficients at scale $j = 3$, where the relation is now represented by dashed black lines e.g. the bottom-right set of 4 dark yellow $j = 3$ coefficients are directly below the bottom-right dark green $j = 2$ coefficient. Notice how the tree (red arrows) need not extend all the way down to the maximum depth.

via a recurrence relation, the details of which we omit here in the interest of space. We refer the interested reader to [Hawkins & Sambridge \(2015\)](#) for details, including a fast algorithm for computing prior ratios of tree structures.

We leave discussion of our value prior $p(\alpha_i | \mathcal{T}_k, k)$ to [subsection 4.2](#), as this is more physically motivated rather than imposed by our parameterisation. Suffice to say here that we use a generalised Gaussian prior on our wavelet coefficient values, with larger wavelets scales (low j) having a more Gaussian prior than smaller scales (high j).

These three priors are independent, so our overall prior $p(\kappa)$ is given by the simple product

$$p(\kappa) = p(\mathcal{T}_k | k) p(k) \prod_{i=1}^k p(\alpha_i | \mathcal{T}_k, k). \quad (10)$$

3.3. Proposals

In this birth-death sampler we have three classes of proposals: (1) value proposals, where the value of a wavelet coefficient is modified; (2) birth proposals, where a new coefficient is added to the tree; and (3) death proposals, where a coefficient is removed from the tree. We divide up our wavelet coefficient space into three sets, $\mathcal{S}_v, \mathcal{S}_b, \mathcal{S}_d$, which contain the coefficients at which a value, birth or death proposal, respectively, can be performed. An example is shown in [Figure 2](#). The value set \mathcal{S}_v corresponds to the current tree, and will always include at least the root of the tree, hence $|\mathcal{S}_v| \geq 1$. The death set \mathcal{S}_d is the “low-hanging fruit” of the tree, i.e. coefficients without any active child coefficients. The birth set

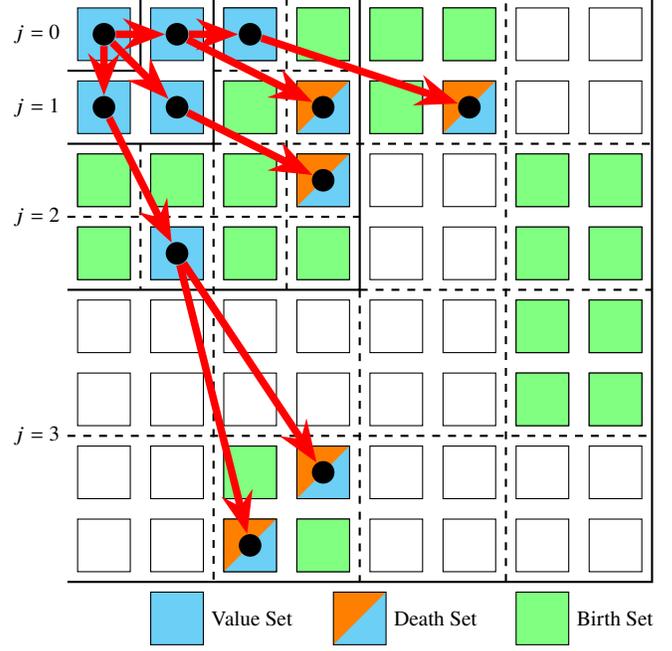


FIG. 2.— Proposal sets for the 2D wavelet tree parameterisation. The wavelet coefficients are arranged as in [Figure 1](#), although here they are coloured by the proposal that may be performed at that coefficient. Value, death and birth proposals can be made at the blue, orange and green coefficients, respectively. Note that the death set is a subset of the value set, and the birth set is disjoint from the other two sets.

\mathcal{S}_b is all the currently inactive child coefficients of the current tree. Notice that the death set is a proper subset of the value set, $\mathcal{S}_d \subset \mathcal{S}_v$, and the birth set is disjoint from the other two sets, $\mathcal{S}_b \cap \mathcal{S}_v = \mathcal{S}_b \cap \mathcal{S}_d = \emptyset$.

The value proposal distribution is given by

$$q_v(\kappa' | \kappa) = \frac{q(\Delta \alpha_i | i)}{|\mathcal{S}_v|}, \quad (11)$$

where the numerator is the distribution from which a perturbation to wavelet coefficient i is drawn. We choose this to be a zero-mean Normal distribution with variance tuned to achieve an acceptance rate of roughly 20–40%. The denominator covers the probability of choosing to perturb the i^{th} coefficient.

Birth proposals are performed by selecting a coefficient i from the birth set \mathcal{S}_b and initialising it with a value drawn from the value prior. In the case where the birth set is empty, i.e. the current tree spans all coefficients, the proposal probability is zero. Hence, we have

$$q_b(\kappa' | \kappa) = \frac{p(\alpha_i | \mathcal{T}_k, k)}{|\mathcal{S}_b|}. \quad (12)$$

Death proposals simply involve choosing a coefficient i to remove from the death set \mathcal{S}_d , so

$$q_d(\kappa' | \kappa) = \frac{1}{|\mathcal{S}_d|}. \quad (13)$$

The acceptance criterion ([Equation 6](#)) requires the reverse proposal distributions for each of these proposals $q(\kappa | \kappa')$. The value proposal distribution is symmetric, so the reverse value proposal is equal to the forward proposal. As for the birth and death proposals, these are the reverse proposals of each other, i.e. the reverse of proposing the birth of coefficient i from \mathcal{S}_b is proposing its death from \mathcal{S}'_d , where the prime here denotes

the set after the proposal. So the reverse proposals are

$$q_v(\kappa|\kappa') = q_v(\kappa'|\kappa), \quad (14)$$

$$q_b(\kappa|\kappa') = q_d(\kappa'|\kappa), \quad (15)$$

$$q_d(\kappa|\kappa') = q_b(\kappa'|\kappa). \quad (16)$$

3.4. Acceptance criteria

Recall the acceptance criterion (Equation 6) depends on the prior, likelihood and proposal ratios, and a Jacobian. The likelihood ratio remains the same regardless of the proposal, and we denote this as $\mathcal{L}(\kappa', \kappa) = p(\gamma|\kappa') / p(\gamma|\kappa)$. The prior and proposal ratios depend on the type of proposal performed.

For value proposals the tree structure does not change, so the prior ratio is simply whether or not the proposed coefficient value is more likely than the current value

$$\frac{p(\kappa')}{p(\kappa)} = \frac{p(\kappa'_i|\mathcal{T}_k, k)}{p(\kappa_i|\mathcal{T}_k, k)}. \quad (17)$$

As for the proposal ratio, the reverse distribution is equal to the forward distribution so the ratio is 1. Since there is no change in dimension of the tree, the Jacobian will always be the identity, with determinant equal to 1. Hence the acceptance criterion for value proposals is

$$\alpha_v(\kappa', \kappa) = \min \left\{ 1, \frac{p(\kappa'_i|\mathcal{T}_k, k)}{p(\kappa_i|\mathcal{T}_k, k)} \mathcal{L}(\kappa', \kappa) \right\}. \quad (18)$$

For birth proposals, the tree structure changes. Using our uniform prior for the number of coefficients in the tree, this part of the prior ratio is unity, $p(k+1) / p(k) = 1$. The prior values also all cancel out except at the proposed new coefficient i . So the overall birth prior ratio is, dropping the conditional dependencies for clarity

$$\frac{p(\kappa')}{p(\kappa)} = \frac{p(\mathcal{T}_{k+1})p(\kappa_i)}{p(\mathcal{T}_k)}. \quad (19)$$

Using reverse distributions above, the proposal ratio is

$$\frac{q(\kappa|\kappa')}{q(\kappa'|\kappa)} = \frac{|\mathcal{S}_b|}{|\mathcal{S}'_d|p(\kappa_i)}. \quad (20)$$

So the birth proposal acceptance criterion is

$$\alpha_b(\kappa', \kappa) = \min \left\{ 1, \frac{p(\mathcal{T}_{k+1})}{p(\mathcal{T}_k)} \frac{|\mathcal{S}_b|}{|\mathcal{S}'_d|} \mathcal{L}(\kappa', \kappa) \right\}. \quad (21)$$

For the Jacobian, it is a simple case of Green (1995)'s dimension matching without any transformation of random variables. We can denote our parameter space θ as a set of k tuples each containing a unique index t_i denoting a position in the tree and the value of the wavelet coefficient at that position κ_i . Dimension matching then requires us to sample random variables u and w for the position and value of the new coefficient, respectively, and then setting the proposal θ' to be some function of u and w . Using the proposals defined previously, that function is simply the identity. For all the pre-existing parameter ($i = 1 \dots k$) there is no change, and for the new coefficient we set $(t_{k+1}, \kappa_{k+1}) = (u, w)$, hence there is no transformation of the random variable and the Jacobian is the identity matrix.

Similar reasoning and considering that the Jacobian of the death proposal is the inverse of the birth proposal Jacobian leads to the death proposal acceptance criterion,

$$\alpha_d(\kappa', \kappa) = \min \left\{ 1, \frac{p(\mathcal{T}_{k-1})}{p(\mathcal{T}_k)} \frac{|\mathcal{S}_d|}{|\mathcal{S}'_b|} \mathcal{L}(\kappa', \kappa) \right\}. \quad (22)$$

4. MASS-MAPPING WITH TRANS-DIMENSIONAL TREES

In this section we give details of how we use the trans-dimensional tree structure described previously for mass-mapping. In particular we describe the forward operator and choice of prior used for the following sections.

4.1. Forward Operator

Our trans-dimensional MCMC samples the wavelet coefficients of a convergence map κ . The convergence map is then $\kappa = W^{-1}\gamma$, where W^{-1} is the inverse wavelet transform. To use the KS kernel in Equation 3, denoted by D we need forward and inverse fast Fourier transforms, F, F^{-1} , respectively. As such, obtaining the shear from sampled convergence wavelet coefficients is given by

$$\gamma = F^{-1}DFW^{-1}\kappa. \quad (23)$$

The input data for our inversions are gridded galaxy shear measurements. Gridding is necessary to allow the galaxy ellipticity due to shear to emerge out of the shape noise due to intrinsic ellipticity. We make the common assumption that the intrinsic ellipticities of galaxies are Gaussian randomly oriented with zero mean, and thus we take the observed shear in a small region of sky (pixel) to be the mean of the ellipticities of the galaxies in that region. However, as a result of the gridding we are left with fewer data points, and likely insufficient data to produce a high-resolution convergence map. Further, depending on the coarseness of the gridding it is possible that some pixels contain no galaxy measurements, as is common in weak lensing surveys, which may lead to signal leakage between Fourier modes.

We assume that the shear measurements in each pixel follow a Gaussian distribution of known variance σ_i^2 . As such we set the likelihood of observing shear measurements γ for a given convergence map wavelet coefficients κ to be

$$p(\gamma|\kappa) = \frac{1}{\sqrt{(2\pi)^N |C_d|}} \exp \left\{ -\frac{(\Phi\kappa - \gamma)^T C_d^{-1} (\Phi\kappa - \gamma)}{2} \right\}, \quad (24)$$

where Φ is the operator defined in Equation 23 and C_d is the data covariance matrix. We typically assume no correlation between neighbouring pixels, and as such C_d is a diagonal matrix with all the σ_i^2 along the diagonal. We also assume that the real and imaginary parts of the σ_i^2 are equal.

4.2. Generalised Gaussian Prior on Wavelet Coefficients

Strictly, sparsity is measured by the ℓ_0 -norm, which counts the non-zero elements in a vector. However, this norm is non-convex, and it has been shown that the convex ℓ_1 -norm, or equivalently, the Laplace distribution, produces comparable results, and so is now widely used for sparsity promotion (Donoho 2006; Candès et al. 2011). However, thanks to the multiresolution nature of the wavelets we can impose different priors for different length scales. In particular, we expect the wavelet coefficients of larger scale structure to be more Gaussian than at smaller scales. Starck et al. (2021) had a similar idea and solved for sparsity-based component to capture non-Gaussian structure and a separate Gaussian random field. Alternatively, McEwen et al. (2017) suggested fitting a generalised-Gaussian distribution (GGD) to each wavelet scale. The probability density function of a GGD is given by

$$p(x|\mu, \sigma, \beta) = \frac{\beta}{2\sigma\Gamma(\beta^{-1})} \exp \left\{ -\left| \frac{x - \mu}{\sigma} \right|^\beta \right\}, \quad (25)$$

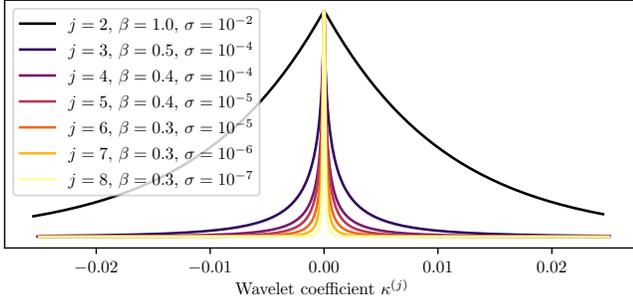


FIG. 3.— Fitted generalised Gaussian distributions of wavelet coefficients for the Bolshoi 7 & 8 N-body galaxy cluster simulations, normalised to have unit height. The scaling function coefficients and the largest wavelet scale coefficients $j = 1$ have been omitted, as there are not enough coefficients for a fit to be truly meaningful. The fitted scale, σ , and shape, β parameters are indicated.

where $\Gamma(\cdot)$ is the gamma function, μ is the mean, σ is the scale parameter (c.f. standard deviation) and β is the shape parameter. For $\beta = 2$ this gives the Gaussian distribution, $\beta = 1$ gives the Laplacian, and as $\beta \rightarrow \infty$ this tends to a uniform distribution in the range $[\mu - \sigma, \mu + \sigma]$. McEwen et al. (2017) suggested learning the scale and shape parameters for each wavelet scale based on simulation images. Figure 3 show the fitted GGDs of wavelet coefficients at each scale for the Bolshoi 7 & 8 galaxy cluster simulation convergence images (Klypin et al. 2011). The fitted scale and shape parameters are also show, and clearly the larger scale wavelet coefficients are more Gaussian than the smaller scale coefficients. We use the distributions in Figure 3 as loose guide for tuning the shape and scale parameters for each wavelet scale for the application to simulation data in the following section. It would be unreasonable to use the exact distributions as they would not be known in the real data case.

In the wavelet tree parameterisation, the root of the tree is a single pixel approximation of the κ map, i.e. the image mean. However, the mass-sheet degeneracy means κ can only be determined up to an unknown constant λ , $\kappa_\lambda = \lambda\kappa + (1 - \lambda)$. We can then set the prior of the tree root to be extremely tight around 0, as it cannot be determined from the data.

5. APPLICATION TO SIMULATIONS

In this section we present the results of tests of our trans-dimensional tree method on datasets generated from simulated cluster-scale convergence maps.

5.1. Data

We begin with a ground truth convergence map κ , extracted from the Bolshoi N-body simulation (Klypin et al. 2011). We use the same simulation catalogues as Lanusse et al. (2016) and Price et al. (2021), extracted using the CosmoSim¹ website. This assumed a redshift of $z = 0.3$, a 10×10 arcmin² field of view, and the convergence maps have been normalised with respect to lensing sources at infinity. The ground truth image, shown in Figure 4

From the ground truth convergence maps we generate noisy synthetic shear data γ from

$$\gamma = F^{-1}DF\kappa + n, \quad (26)$$

where n is drawn from a zero-mean Gaussian distribution $\mathcal{N}(0, \sigma^2)$, σ^2 is the data covariance determined by the number

of observed galaxies in a pixel and an assumed intrinsic galaxy ellipticity dispersion of 0.37 (Price et al. 2021),

$$\sigma^2 = \frac{0.37^2}{\sqrt{2N}} \quad (27)$$

where N is the expected number of galaxies in a pixel for a given number density of galaxies per arcmin², n_{gal} . For example, *Euclid* is expected to be able to see about $n_{\text{gal}} \sim 30$ galaxies per arcmin² (e.g Laureijs et al. 2011). To make the data somewhat more realistic, we apply a random mask to the noisy shear data whereby we mask out 1% of pixels to simulate regions with no galaxy observations.

5.2. Method

We invert the synthetic shear data for a convergence map using the trans-dimensional tree method detailed in Section 3. The wavelets we use are the Cohen-Debauchies-Feauveau 9/7 wavelet family (Cohen et al. 1992), as these produced the best results for Hawkins & Sambrige (2015). We run a single Markov chain for at least 10^6 steps and check that the likelihood and number of model parameters k has converged. If needed we restart the chain from where it ended and extend the chain for as long as necessary. The generalised Gaussian prior parameters (shape and scale) are tuned for each wavelet scale, except the lowest scale which represents the mass-sheet degeneracy, so as to achieve an appropriate acceptance rate, typically around 20–50%. As in Price et al. (2021) we evaluate the similarity between our convergence solutions (mean or best-fitting) and the ground truth using two quantities: the signal-to-noise ratio (SNR) in dB to assess the overall difference

$$\text{SNR}(\kappa^*) = 10 \times \log_{10} \left(\frac{\|\kappa\|_2^2}{\|\kappa - \kappa^*\|_2^2} \right) \quad (28)$$

and the Pearson correlation coefficient (r) to assess the structural similarities.

$$r(\kappa^*) = \frac{\sum_{i=1}^M (\kappa_i - \langle \kappa \rangle)(\kappa_i^* - \langle \kappa^* \rangle)}{\sqrt{\sum_{i=1}^M (\kappa_i - \langle \kappa \rangle)^2} \sqrt{\sum_{i=1}^M (\kappa_i^* - \langle \kappa^* \rangle)^2}} \quad (29)$$

Here κ^* denotes our solution convergence map and $\langle \cdot \rangle$ denotes an average over the M pixels.

5.3. Results

5.3.1. Recovery of a clean high-resolution mass-map

Figure 4 shows the ground truth map, the mean and highest posterior (MAP estimate) trans-dimensional MCMC point solutions and the range of the highest posterior density region at the 99% credible interval level. This is a very clean data example. Both the mean and highest posterior point estimates recover the three main high convergence regions well, although some of the fainter sources are missing, possibly lost in the noise or smoothed out when averaging over many samples in the case of the mean solution. The highest posterior point estimate recovers nicely the peak at the core of the lowermost cluster, which seems to have been lost in the mean solution. In terms of the overall SNR and correlation, however, the mean solution is better than the MAP. This is somewhat to be expected, as it is the nature of MCMC that the true MAP solution is unlikely to be found, and individual estimates are likely to have a few artefacts.

¹ <http://www.cosmosim.org>

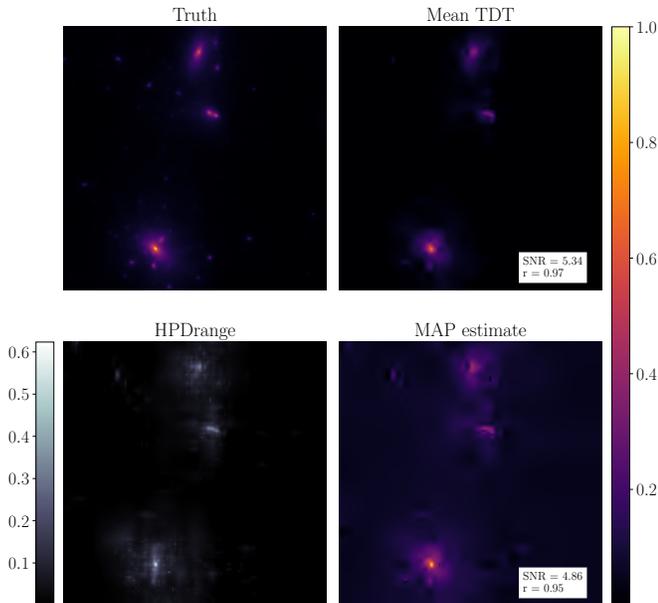


FIG. 4.— (top left) Ground truth convergence field from which synthetic shear measurements were obtained. (top right) Mean solution from trans-dimensional MCMC. (bottom right) Highest posterior sample (estimate of the MAP solution) from trans-dimensional MCMC. (bottom left) Size of the highest posterior density region as a measure of uncertainty.

The quantified uncertainty is represented by the 99% credible interval range, symmetric about the peak of the histogram of pixel values. Assuming a unimodal histogram, this is the highest posterior density region. The lateral extent of the three main clusters in the uncertainty map are slightly larger than those in the point solutions. As well as highlighting the three main peaks, where the detailed structures are not quite resolved, the uncertainty map picks out some of the smaller, fainter structures that have been missed in the point solutions. For example, a faint source in the top left quadrant is identified, as are faint structures in the vicinity of topmost and lowermost main peaks.

Figure 5 shows a histogram of the number of active tree nodes (non-zero wavelet coefficients, k) for all the samples in the trans-dimensional MCMC chain, as well as the evolution of the likelihood $p(\gamma)|p(\kappa)$ and the number of wavelet coefficients k as the chain progresses. The number of coefficients generally increases, as designed by the parameterisation. As the number of coefficients increases, the likelihood expectedly improves. However, the number of parameters does not increase indefinitely and converges at around 250 coefficients. In this case, the ground truth is sampled on a 256×256 pixel grid, meaning the maximum potential size of the parameter space is $k_{\max} = 65536$. This is a massive reduction of the parameter space.

5.3.2. Low-resolution and high-noise inversions

In the example shown previously, the noise level was kept relatively low, using $n_{gal} = 5000$ galaxies per arcmin² which corresponds to just under 8 galaxies per pixel. The *Euclid* telescope is expecting to observe around 30 galaxies per arcmin² of sky (Laureijs et al. 2011). For this 10×10 arcmin patch of sky, that would correspond to on average less than 0.05 galaxies per pixel. During our initial experiments, we found this noise level to be too high for our method to work. Adding or removing parameters would cause so little a difference to the likelihood (data fit) that almost any proposal would be ac-

cepted, and since births and deaths are proposed in the same proportions, the tree would never grow.

The solution that we found to beat down the noise is to decrease the resolution. At 32×32 pixels, a *Euclid*-like noise level would correspond to about 3 galaxies per pixel for this patch of sky. Figure 6 shows the results (mean and uncertainty) of our trans-dimensional inversions at this lower resolution for varying noise levels. Also shown are the KS solution and the KS solution with optimum smoothing. The smoothing uses a simple 2D Gaussian kernel, the optimum size of which was determined by a simple linear search of the kernel size σ and identifying which kernel produced the best SNR. This would obviously not be possible in practice with real data, as it requires knowledge of the true convergence field. As such, this optimum solution represents the top-end of what could be possible with KS. At all noise levels, trans-dimensional MCMC performs better than the optimally smoothed KS solution in terms of SNR and correlation with the ground truth. Visually, even at the highest noise levels, the lowermost peak is more tightly constrained by MCMC. The additional uncertainty information obtained from MCMC provides further constraints on the locations of peaks, although, somewhat surprisingly, the uncertainty is more laterally spread in the low noise cases than in the high noise cases. The size of the uncertainty does decrease as the noise as noise decreases, as expected.

6. DISCUSSION

The trans-dimensional MCMC algorithm is able to significantly decrease the size of the parameter space, searching only a few hundred parameters rather than tens of thousands in the high resolution case (Figure 4). The uncertainties also provide further constraints on the lateral extent of the peaks of the convergence field. All our results also compare favourably to the standard KS method, both in terms of the overall fidelity of the reconstruction and in its ability to fully quantify uncertainties (Figure 6). Importantly, our results compare favourably to the optimally smoothed KS solution, which is unobtainable in practice. Furthermore, this sampling problem is far from computationally expensive, thanks to the reduction in parameter space. All the inversions performed were conducted on a 2020 MacBook Pro with an Apple M1 processor, with the high resolution example taking only 5 hours.

With a new method that more accurately recovers mass maps from noisy shear data and also quantifies full uncertainty, this opens up an exciting prospect of new results relating to dark matter distributions in galaxy clusters. This method could be applied to high-resolution shear catalogues from the Hubble Frontier Fields (Koekemoer et al. 2014) or new data from *Euclid* (Laureijs et al. 2011) for new mass-maps and comparisons with, for example, data from the Chandra X-ray Observatory (Weisskopf et al. 2000) to investigate the collisional nature of dark matter, as in Harvey et al. (2015). Harvey et al. (2015) used a parametric model to reconstruct the density distribution of their colliding galaxy systems (Navarro et al. 1997; Jullo et al. 2007). Our non-parametric method will thus relax some of the physical assumptions made in their density reconstructions, providing additional constraints and validation of their findings.

The main limitations of this trans-dimensional sampler are two-fold. Firstly, convergence to the peak of the posterior is difficult to determine. In some cases, the sampler may seem to have converged, then a new tree node is introduced that causes a large jump in posterior probability, bringing the sampler into a new region of parameter space to be explored. This is a

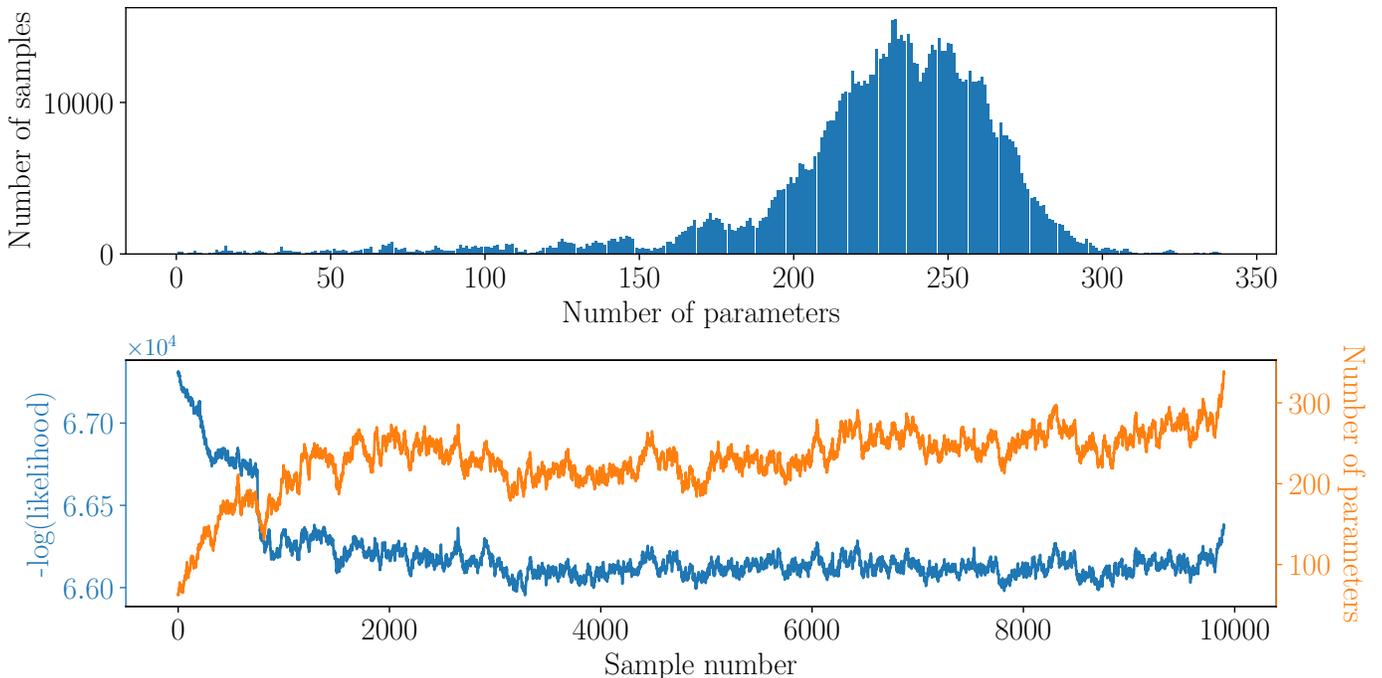


FIG. 5.— History of the MCMC chain that produces the results in Figure 4. (top) Histogram of number of wavelet coefficients. (bottom) Evolution over the Markov chain of the likelihood (blue) and number of wavelet coefficients (orange). Note that the samples in the bottom plot have been thinned, i.e. only every 100th sample is shown, but have not been thinned in the top plot, i.e. every sample is shown. This explains why the top plot contains more samples (area under the histogram) than the bottom plot (length of x axis).

general problem with MCMC methods. While there is some theoretical understanding about the convergence properties of various probabilistic methods (e.g. Roberts & Smith 1994; Mengersen & Tweedie 1996; Pereyra 2016; Durmus et al. 2017), how this translates into number of iterations is not straightforward, and so in practice it is often the case of simply continuing for as long as possible. A common choice is also to run multiple chains in parallel, starting each chain in a different region of parameter space and converging to the peak of the posterior. In a standard trans-dimensional MCMC this could be desirable, as one could start with very different models each with a different number of parameters. For this wavelet tree parameterisation on the other hand, the initial parameter space only has one dimension, by design. Thus the multiple chains would be starting in very similar points of parameter space, negating the benefit of having parallel chains. This could be alleviated with some Parallel Tempering techniques (Swendsen & Wang 1986; Falcioni & Deem 1999; Brooks et al. 2011; Sambridge 2013), whereby each chain samples a slightly modified posterior and there is information exchanged between chains. For example, each chain will sample the distribution

$$\pi(\kappa|T) = p(\kappa|\gamma)^{1/T}, \quad (30)$$

for a given temperature $T > 0$. For $T = 1$ this is equal to the desired posterior, and for $T > 1$ this is a slightly smoother version of the posterior. With information being exchanged between chains, there should, in principle, be better mixing of the MCMC samples (Sambridge 2013).

The second main limitation is the handling of very noisy data. As previously mentioned, when the noise level gets too high the addition or removal of parameters makes little difference to the posterior probability. Thus any proposal is generally accepted and the tree struggles to grow. We found this to also be an issue with the smallest wavelet scales, where the addition of a very small wavelet would barely affect the

likelihood. This resulted in anomalously high birth and death acceptance rates among the smallest wavelet scales. However, there is no tuning parameter (scale and shape of the prior distributions) that has an affect on the acceptance criteria of birth and death proposals (see Equations 21 and 22). This is due to the choice of drawing the value of new coefficients from the prior. While other choices are of course possible, none are physically well-motivated. As a result it is quite difficult to resolve the smallest and faintest structures, particularly as they get hidden by noisy data.

The wavelet tree parameterisation used here is not the only parameterisation that can be used for a trans-dimensional MCMC. Indeed, a popular choice in many geophysical applications is to use Voronoi cells (Voronoi 1908; Bodin et al. 2012; Young et al. 2013; Dettmer et al. 2014; Zhang et al. 2018), where the image space is divided into k cells defined by a node location such that all the points in a given cell are closer to their cell node than any other cell node. The value of the field for which one is inverting is generally taken to be constant in each cell. The nodal density of the cells tends to increase near regions that have the most influence on the posterior, e.g. where there is more data or strongly heterogenous regions. This could conceivably be an excellent option for the mass-mapping problem where the data coverage is dictated by the positions of galaxies on the sky, and gridding may cause holes in the observed shear field. However, the Voronoi cells may not necessarily form a sparse basis, which may make it difficult to preserve non-Gaussian structures in the resultant convergence map.

7. CONCLUSIONS

In this work we have used, for the first time, a trans-dimensional MCMC sampler to build cosmological mass-maps in a probabilistic manner, with a sparsity-promoting wavelet prior. The wavelet tree parameterisation is designed so that the large scale information is resolved first, with smaller

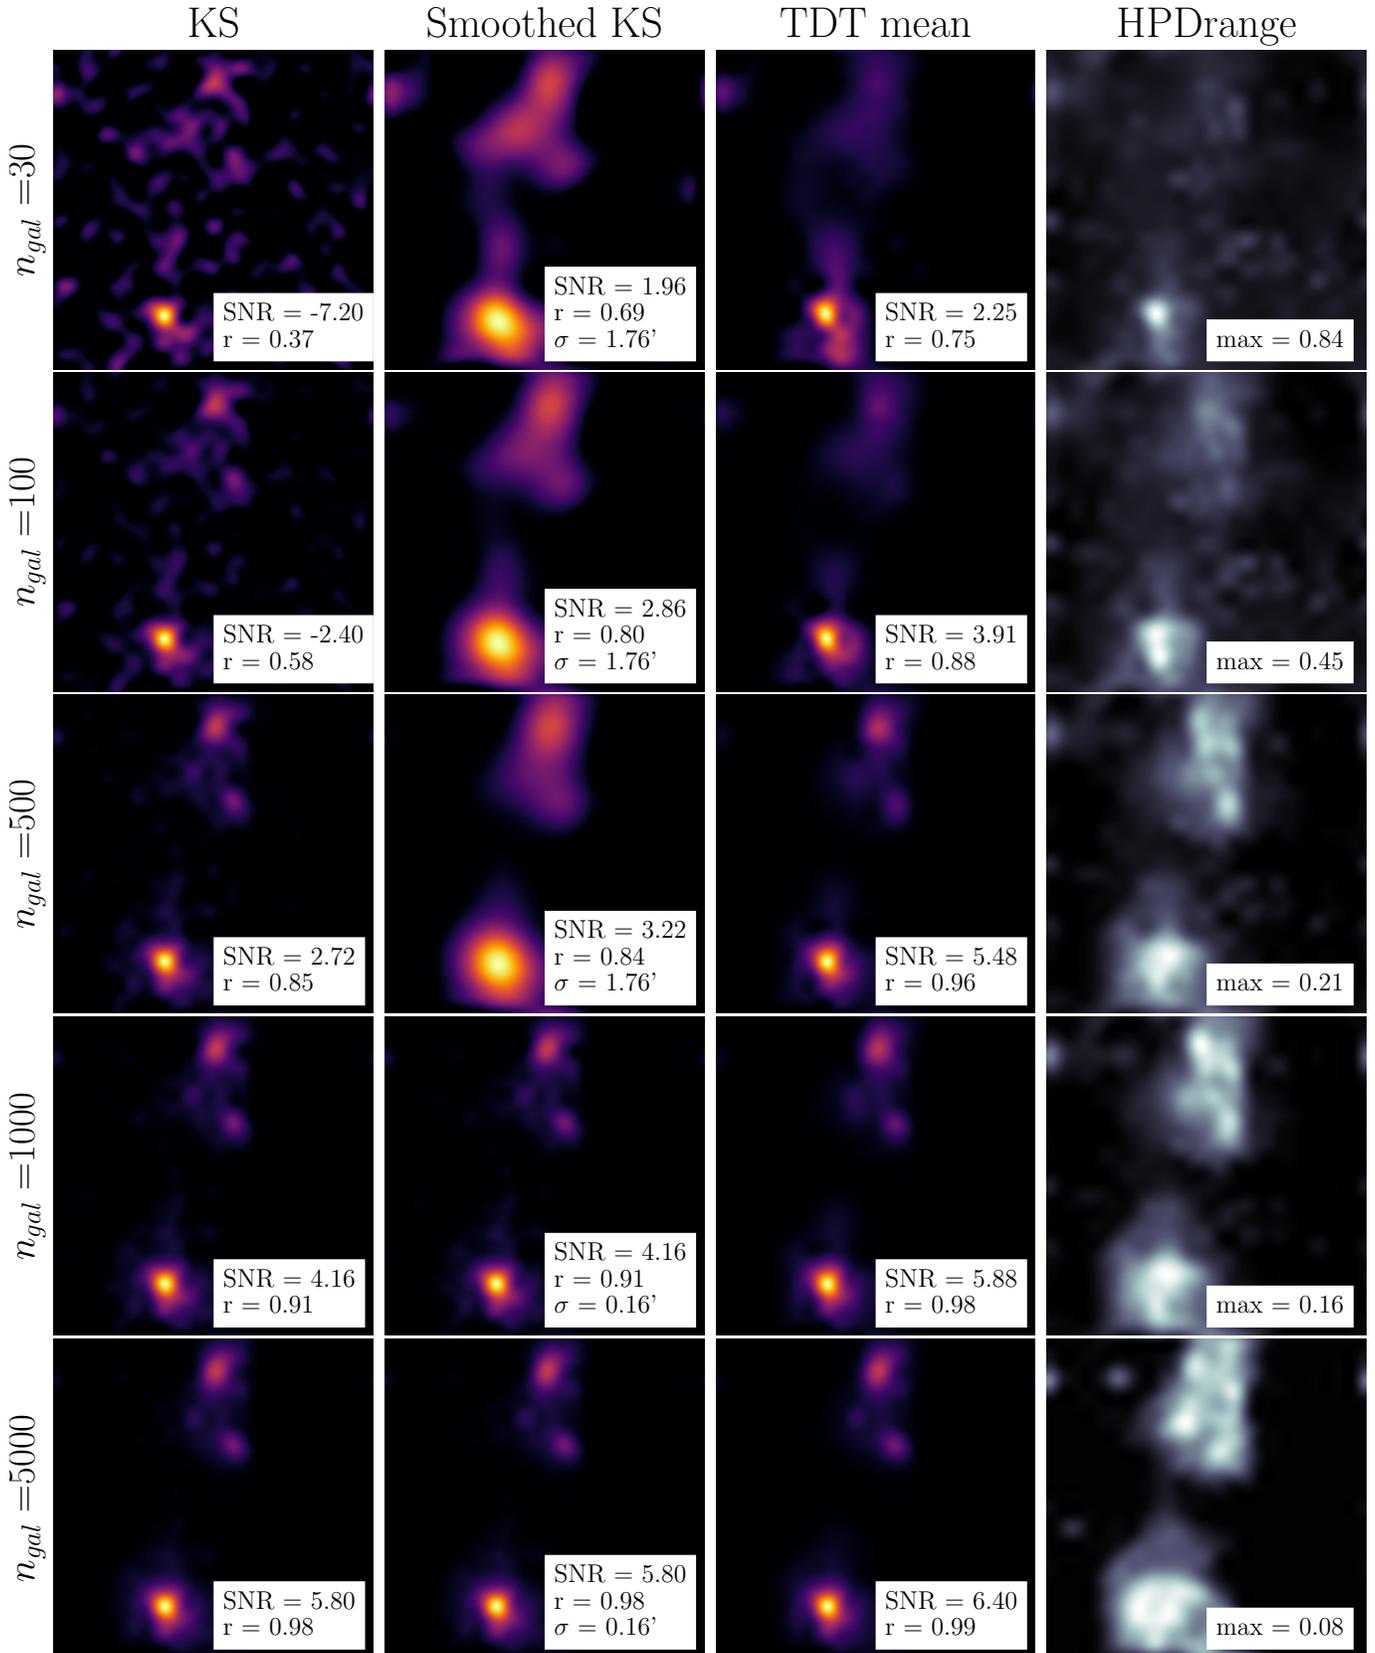


FIG. 6.— Comparison of trans-dimensional trees against KS at various levels of noise. The ground truth image is the same as that in Figure 4, resampled onto a 32×32 grid to beat down the noise by increasing the number of average galaxies per pixel. Each row is for a given noise level, defined by an average number of galaxies per arcmin², n_{gal} . Noise decreases from top to bottom. Each column shows, from left to right, the KS reconstruction, the KS reconstruction with optimum smoothing, the mean of the trans-dimensional MCMC, and the uncertainty produced from MCMC. The quality of the reconstructions in the first three columns is indicated by the SNR and correlation coefficient. The optimum Gaussian smoothing kernel size σ in arcmin is also indicated (second column), as is maximum uncertainty value (last column).

scale detail gradually added as required by the data. Instead of using the standard Laplacian distribution as the sparsity-promoting prior, we exploit the multi-resolution nature of the wavelets and the generalised Gaussian distribution to create a scale-dependent prior that is more sparse for the smaller wavelets. This approach successfully recovers mass maps from simulated data better than the standard KS method at high noise levels, including levels expected for upcoming surveys, opening up the possibility for new high-resolution mass maps and inferences about the nature of dark matter.

ACKNOWLEDGEMENTS

A.M. is supported by the STFC UCL Centre for Doctoral Training in Data Intensive Science (grant number

ST/P006736/1). The authors are grateful to Rhys Hawkins and Malcolm Sambridge for making their original code publicly available.

DATA AVAILABILITY

Code for this work is made available at <https://github.com/auggiemarignier/tdtmassmapping>. It makes use of and adapts the codes made available by Rhys Hawkins at <https://github.com/rhyshawkins/TDTbase> and <https://github.com/rhyshawkins/TDTWavetomo2D>. The weak lensing simulation data used is available at <http://www.cosmosim.org/>.

REFERENCES

- Alsing J., Heavens A., Jaffe A. H., Kiessling A., Wandelt B., Hoffmann T., 2015, *Monthly Notices of the Royal Astronomical Society*, 455, 4452
- Bartelmann M., Schneider P., 2001, *Physics Reports*, 340, 291
- Bodin T., Sambridge M., 2009, *Geophysical Journal International*, 178, 1411
- Bodin T., Sambridge M., Gallagher K., Rawlinson N., 2012, *Journal of Geophysical Research: Solid Earth*, 117
- Brooks S., Gelman A., Jones G., Meng X.-L., 2011, *Handbook of markov chain monte carlo*. CRC press
- Burdick S., Waszek L., Lekić V., 2019, *Earth and Planetary Science Letters*, 528, 115789
- Candès E. J., Eldar Y. C., Needell D., Randall P., 2011, *Applied and Computational Harmonic Analysis*, 31, 59
- Chang C., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 3165
- Clowe D., Bradač M., Gonzalez A. H., Markevitch M., Randall S. W., Jones C., Zaritsky D., 2006, *The Astrophysical Journal*, 648, L109
- Cohen A., Daubechies I., Feauveau J.-C., 1992, *Communications on pure and applied mathematics*, 45, 485
- Cornish N. J., Littenberg T. B., 2007, *Physical Review D*, 76, 083006
- Daubechies I., 1992, *Ten lectures on wavelets*. SIAM
- Dettmer J., Benavente R., Cummins P. R., Sambridge M., 2014, *Geophysical Journal International*, 199, 735
- Dodelson S., 2017, *Gravitational Lensing*. Cambridge University Press, Cambridge, doi:10.1017/9781316424254
- Donoho D., 2006, *IEEE Transactions on Information Theory*, 52, 1289
- Durmus A., Moulines E., Saksman E., 2017, *arXiv*
- Falcioni M., Deem M. W., 1999, *The Journal of chemical physics*, 110, 1754
- Feder R., Daylan T., 2018, in *American Astronomical Society Meeting Abstracts #231*, p. 151.04
- Fiedorowicz P., Rozo E., Boruah S. S., Chang C., Gatti M., 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 73
- Gallagher K., 2012, *Journal of Geophysical Research: Solid Earth*, 117
- Geyer C. J., Møller J., 1994, *Scandinavian journal of statistics*, pp 359–373
- Green P. J., 1995, *Biometrika*, 82, 711
- Harvey D., Massey R., Kitching T., Taylor A., Tittley E., 2015, *Science*, 347, 1462
- Hawkins R., Sambridge M., 2015, *Geophysical Journal International*, 203, 972
- Heavens A., 2009, *Nuclear Physics B - Proceedings Supplements*, 194, 76
- Jeffrey N., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 2871
- Jeffrey N., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 4626
- Jullo E., Kneib J.-P., Limousin M., Eliasdottir A., Marshall P., Verdugo T., 2007, *New Journal of Physics*, 9, 447
- Kaiser N., Squires G., 1993, *The Astrophysical Journal*, 404, 441
- Karnesis N., et al., 2014, *Physical Review D*, 89, 062001
- Klypin A. A., Trujillo-Gomez S., Primack J., 2011, *The Astrophysical Journal*, 740, 102
- Koekemoer A. M., et al., 2014, in *American Astronomical Society Meeting Abstracts #223*, p. 254.02
- Lanusse F., Starck J. L., Leonard A., Pires S., 2016, *Astronomy & Astrophysics*, 591
- Laureijs R., et al., 2011, *arXiv preprint arXiv:1110.3193*
- Malinverno A., 2002, *Geophysical Journal International*, 151, 675
- Mallat S., 1989, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 674
- Marignier A., McEwen J. D., Ferreira A. M. G., Kitching T. D., 2022, *Royal Astronomical Society Techniques and Instruments (in press)*
- McEwen J. D., Feeny S. M., Peiris H. V., Wiaux Y., Ringeval C., Bouchet F. R., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 4081
- Mengersen K. L., Tweedie R. L., 1996, *The Annals of Statistics*, 24, 101
- Minsley B. J., 2011, *Geophysical Journal International*, 187, 252
- Munshi D., Coles P., 2017, *Journal of Cosmology and Astroparticle Physics*, 2017, 010
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *The Astrophysical Journal*, 490, 493
- Neal R. M., 2012, in *Handbook of Markov Chain Monte Carlo*. CRC Press, pp 113–162
- Peebles P. J. E., 1980, *The large-scale structure of the universe*
- Pereyra M., 2016, *Statistics and Computing*, 26, 745
- Pereyra M., 2017, *SIAM Journal on Imaging Sciences*, 10, 285
- Piana Agostinetti N., Giacomuzzi G., Malinverno A., 2015, *Geophysical Journal International*, 201, 1598
- Porqueres N., Heavens A., Mortlock D., Lavaux G., 2021, *Monthly Notices of the Royal Astronomical Society*, 502, 3035
- Price M. A., Cai X., McEwen J. D., Pereyra M., Kitching T. D., 2020a, *Monthly Notices of the Royal Astronomical Society*, 492, 394
- Price M. A., McEwen J. D., Pratley L., Kitching T. D., 2020b, *Monthly Notices of the Royal Astronomical Society*, 500, 5436
- Price M. A., McEwen J. D., Cai X., Kitching T. D., Wallis C. G. R., 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 3678
- Remy B., Lanusse F., Jeffrey N., Liu J., Starck J.-L., Osato K., Schrabback T., 2022, *arXiv*
- Roberts G., Smith A., 1994, *Stochastic Processes and their Applications*, 49, 207
- Sambridge M., 2013, *Geophysical Journal International*, 196, 357
- Starck J. L., Themelis K. E., Jeffrey N., Peel A., Lanusse F., 2021, *Astronomy & Astrophysics*, 649
- Swendsen R. H., Wang J.-S., 1986, *Physical review letters*, 57, 2607
- Tkalčić H., Young M., Bodin T., Ngo S., Sambridge M., 2013, *Nature Geoscience*, 6, 497
- Van Waerbeke L., et al., 2013, *Monthly Notices of the Royal Astronomical Society*, 433, 3373
- Voronoi G., 1908, *J. reine angew. Math.*, 134, 198
- Wallis C. G. R., Price M. A., McEwen J. D., Kitching T. D., Leistedt B., Plouviez A., 2021, *Monthly Notices of the Royal Astronomical Society*
- Weisskopf M. C., Tananbaum H. D., Van Speybroeck L. P., O'Dell S. L., 2000, *X-Ray Optics, Instruments, and Missions III*, 4012, 2
- Young M. K., Rawlinson N., Bodin T., 2013, *Geophysics*, 78, WB49
- Zhang X., Curtis A., Galetti E., de Ridder S., 2018, *Geophysical Journal International*, 215, 1644

arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.

This paper was built using the Open Journal of Astrophysics \LaTeX template. The OJA is a journal which provides fast and easy peer review for new papers in the astro-ph section of the