

Generative models of astrophysical fields with scattering transforms on the sphere

L. Mousset[1]⋆, E. Allys[1], M. A. Price[2], J. Aumont[3], J.M. Delouis[4], L. Montier[3], and J. D. McEwen[2]

[1] Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, 75005 Paris, France
[2] Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK
[3] Institut de Recherches en Astrophysique et Planétologie, Université de Toulouse, CNRS, CNES, UPS, Toulouse, France
[4] Laboratoire d'Océanographie Physique et Spatiale, Univ. Brest, CNRS, Ifremer, IRD, Brest, France

Received xxx; accepted xxx

ABSTRACT

Scattering transforms are a new type of summary statistics recently developed for the study of highly non-Gaussian processes, which have been shown to be very promising for astrophysical studies. In particular, they allow one to build generative models of complex non-linear fields from a limited amount of data, and have also been used as the basis of new statistical component separation algorithms. In the context of upcoming cosmological surveys, such as LiteBIRD for the cosmic microwave background polarization or Rubin-LSST and Euclid for study of the large scale structures of the Universe, the extension of these tools to spherical data is necessary. We develop scattering transforms on the sphere and focus on the construction of maximum-entropy generative models of several astrophysical fields. We construct, from a single target field, generative models of homogeneous astrophysical and cosmological fields, whose samples are quantitatively compared to the target fields using common statistics (power spectrum, pixel probability density function and Minkowski functionals). Our sampled fields agree well with the target fields, both statistically and visually. These generative models therefore open up a wide range of new applications for future astrophysical and cosmological studies; particularly those for which very little simulated data is available. We make our code available to the community so that this work can be easily reproduced and developed further.

Key words. Physical data and processes, Methods: data analysis, Methods: statistical, Cosmology: Large-scale structure of Universe

1. Introduction

Scattering transforms are a recently developed class of summary statistics for the study of non-Gaussian processes (Mallat 2012; Bruna & Mallat 2013). These statistics, which are built from successive wavelet convolutions and pointwise non-linearities such as a modulus, are inspired by neural networks but do not require any training step. Introduced recently in astrophysics (Allys et al. 2019, 2020), scattering transforms have since demonstrated their ability to characterize highly non-Gaussian processes, for instance for parameter estimation and classification tasks in fields as varied as the interstellar medium (Regaldo-Saint Blancard et al. 2020; Saydjari et al. 2021; Lei & Clark 2023), the large scales structures of the Universe (Allys et al. 2020; Cheng et al. 2020; Cheng & Ménard 2021; Valogiannis & Dvorkin 2022a,b), or the epoch of reionization (Greig et al. 2022; Hothi et al. 2024).

Another feature of scattering transforms is that they allow one to build very efficient generative models of physical fields, in a maximum entropy framework (Bruna & Mallat 2019). This allows one to sample new approximate realisations of a given process relying only on its scattering transform statistics, which can be estimated from a small amount of data, sometime even a single example image (Allys et al. 2020; Régaldo-Saint Blancard et al. 2023; Price et al. 2023; Cheng et al. 2024b). One application of such generative models is for training data augmentation for machine learning applications. For instance, it has been shown in Jeffrey et al. (2022), with simulated data, that such scattering transform models constructed from a single polarized microwave dust foreground patch can be sufficient to separate primordial CMB B-modes from this dust emission, even in an artificially challenging mono-frequency approach. Moreover, the framework on which these generative models have been constructed has led to the development of new statistical component separation algorithms, which have for instance been successfully applied to astrophysical data (Regaldo-Saint Blancard et al. 2021; Delouis et al. 2022; Auclair et al. 2024) and seismic signals (Siahkoohi et al. 2023a,b).

While these promising scattering transform generative models have mainly been developed for 2D planar data, the adaptation of these tools to spherical data is necessary for cosmological analysis, especially for the next generation of full sky surveys, such as LiteBIRD (LiteBIRD Collaboration et al. 2023), Euclid (Laureijs et al. 2011) or Rubin-LSST (LSST Science Collaboration et al. 2009). The adaptation of a first generation of scattering transforms to spherical signals was already introduced in McEwen et al. (2022), and used as a form of dimensionality reduction for other machine learning purposes. In this paper, we extend state-of-the-art scattering transforms (Morel et al. 2023; Cheng et al. 2024b), named scattering covariances and abbreviated by SC in the following, to spherical fields.

The extension of scattering transforms to spherical data raises some difficulties: the definition of a directional spherical convolution with oriented filters such as wavelets (McEwen et al. 2015b, 2018); and the transposition of the planar translations which appear in certain scattering transform representa-

⋆ louise.mousset@phys.ens.fr

tions. As a first step, we restrict ourselves to fields with statistically homogeneous fields with properties that do not depend on the position on the sphere. The generalisation beyond statistical homogeneity will be presented in a forthcoming paper. This naturally leads us to cosmological fields, such as a weak lensing field from the Large Scale Structures (LSS) of the Universe, and maps of the thermal Sunyaev-Zeldovitch (tSZ) effect of the Cosmic Microwave Background (CMB). We also consider a map of the Venus surface. In this paper, for all these spherical data, scattering transform generative models are constructed and validated from one single full-sky image.

In Sec. 2, we present the extension of the SC statistics to spherical fields. Then, in Sec. 3, we present SC-based generative models, and discuss their numerical implementation. Finally, in Sec. 4, we present the results obtained with these models for the four non-Gaussian spherical fields studied. Conclusions are presented in Sec. 5.

2. Scattering covariance on the sphere

The Scattering Covariances (SC), or scattering spectra, were previously introduced in Morel et al. (2023) for one dimensional data and in Cheng et al. (2024b) for two dimension planar maps. In this paper we extend those statistics to spherical maps. This section introduces sampling schemes, directional convolutions, and wavelet transforms on the sphere, after which we define the SC statistics.

2.1. Sampling of spherical maps

A spherical field can be represented by its spherical harmonic transform which is the spherical equivalent of the Fourier transform for planar maps. In the following, we work with the usual spherical coordinates $\boldsymbol{\omega} = (\theta, \varphi)$, with colatitude θ and longitude φ. With these coordinates, the spherical harmonic coefficients $I_{\ell m}$ of a spherical field $I(\boldsymbol{\omega})$ defined over the sphere $\mathbb{S}^2$ correspond to the projection onto the spherical harmonics $Y_{\ell m}(\boldsymbol{\omega})$:

$$I_{\ell m} = \int_{\mathbb{S}^2} I(\boldsymbol{\omega}) Y^*_{\ell m}(\boldsymbol{\omega}) \mathrm{d}\Omega(\boldsymbol{\omega}) \,, \tag{1}$$

where $\mathrm{d}\Omega(\boldsymbol{\omega}) = \sin\theta \mathrm{d}\theta \mathrm{d}\varphi$ is the solid angle element. The field can then be represented by its harmonic expansion, given by

$$I(\boldsymbol{\omega}) = \sum_{\ell=0}^{L-1} \sum_{m=-\ell}^{\ell} I_{\ell m} Y_{\ell m}(\boldsymbol{\omega}) \,. \tag{2}$$

The ℓ index is called the multipole and is inversely proportional to angular scales on the sky, while the order m at a given ℓ goes from $-\ell$ to ℓ and captures the anisotropic component of $I(\boldsymbol{\omega})$. The maximum value of ℓ considered, which is $\ell_{\max} = L - 1$, determines the smallest scale in the transform. For a real field, the coefficients satisfy the relation

$$I_{\ell m} = (-1)^m I^*_{\ell -m} \,. \tag{3}$$

The numerical computation of the forward and inverse spherical harmonic transform depends on the sampling in pixel space of the spherical map, for which different choices can be made. For example, in cosmology, the community often adopts Healpix sampling (Górski et al. 2005), in which all pixels have the same area, which can be an advantage in practice. With this sampling the map resolution is given by the `nside` parameter, the number of pixels being equal to $12 \times \texttt{nside}^2$. However, with this sampling, the spherical harmonic transform (as well as the Wigner transform defined below) is not accurate and needs to be refined iteratively. An alternative is instead to use an equiangular sampling, *e.g.* defined in McEwen & Wiaux (2011), for which these transforms can be computed exactly (to machine precision). With this sampling, abbreviated by MW, the angular dimensions of all pixels are the same, and maps are stored as 2D arrays of shape $(N_\theta, N_\varphi) = (L, 2L-1)$. In this paper, the SC statistics computations support various sampling schemes, including Healpix, MW and others, while internal calculations typically adopt sampling schemes, e.g. MW, that afford exact spherical transforms for improved accuracy.

Another operation required, which is also sampling dependent, is the average on the sphere, defined as:

$$\langle I(\boldsymbol{\omega}) \rangle = \frac{1}{4\pi} \int_{\mathbb{S}^2} I(\boldsymbol{\omega}) \mathrm{d}\Omega(\boldsymbol{\omega}). \tag{4}$$

In pixel space, this computation corresponds to

$$\langle I(\boldsymbol{\omega}) \rangle = \frac{1}{4\pi} \sum_p I(\boldsymbol{\omega}_p) \delta\Omega_p, \tag{5}$$

where the sum is done over all pixels p, whose angular positions are noted $\boldsymbol{\omega}_p$. For approximate quadrature $\delta\Omega_p$ can simply represent solid angle at pixel p. Alternatively, some sampling schemes exhibit exact quadrature (McEwen & Wiaux 2011) in which case $\delta\Omega_p$ denotes quadrature weights. When $I_{\ell m}$ has been computed, one directly has

$$\langle I(\boldsymbol{\omega}) \rangle = \frac{1}{2\sqrt{\pi}} I_{00} \,. \tag{6}$$

2.2. Wavelet transform on the sphere

SC statistics are computed from wavelet transforms, which are obtained by convolving an initial map with a set of wavelet filters, where each filter extracts the local information at a particular scale and direction. Wavelet filters need to be localized both in pixel and harmonic space. In this work, we follow McEwen et al. (2015b, 2018) and define the wavelets in harmonic space in separated form

$$\Psi^j_{\ell m} = \sqrt{\frac{2\ell+1}{8\pi^2}} \kappa^j_\ell \zeta_{\ell m} \,, \tag{7}$$

in order to control their angular and directional localisation properties separately, respectively by kernel κ^j_ℓ, with filter scale j, and directional component $\zeta_{\ell m}$. For the explicit definition of these two functions, one can refer to, e.g., McEwen et al. (2015b). The wavelets are designed to satisfy excellent spatial localisation and asymptotic uncorrelation properties (McEwen et al. 2018).

Fig. 1, left panel, shows the $\Psi_{\ell m}$ coefficients of one filter at a specific scale j and, right panel, a cut at $m = 0$ of the full filter set. Note that with our convention, j scales are ordered with ℓ multipoles, meaning that when j increases, the corresponding angular scale decreases (i.e. ℓ increases). Filters are maximum at $\ell \simeq \eta^j$, with support within $\ell \in [\eta^{(j-1)}, \eta^{(j+1)}]$, where η defines the wavelet dilation parameter. In this paper we use dyadic scaling, corresponding to $\eta = 2$. In the following, when performing a convolution with a filter set, the range of scales probed by the wavelets is given by $J_{\min} \leq j \leq J_{\max}$ where $J_{\min} \geq 0$ and $J_{\max} = \mathrm{ceil}\left(\frac{\log(L-1)}{\log\eta}\right)$. The number of scales is given by $J = J_{\max} - J_{\min} + 1$. The angular resolution of the wavelets is

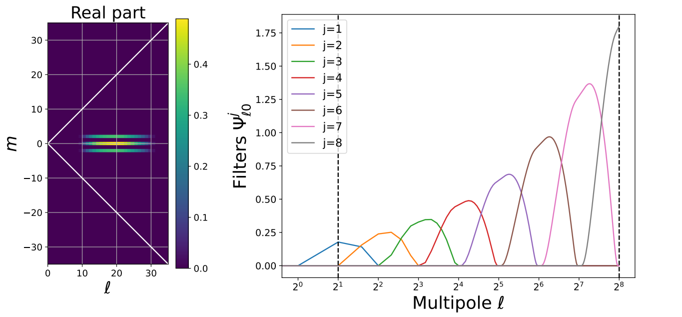

Fig. 1. Left: real part of the $\Psi^{j=4}_{\ell m}$ filter. Right: a cut at $m = 0$ of the full filter set for j spanning from $J_{\min} = 1$ to $J_{\max} = 8$.

parameterized by an integer N, allowing for the sample of $2N-1$ directions (see below)[1]

Wavelet transforms are computed from convolutions between the field under study and this set of wavelets. Various convolutions on the sphere can be considered (Roddy & McEwen 2021). In this work we follow the standard directional convolution formalism presented in, e.g., McEwen et al. (2015b). These convolutions produce a set of spherical maps filtered at different scales (labelled by j) and orientations (labelled by γ).

The directional convolution $I \star \Psi^j$ of a field I with a wavelet Ψ^j consists in applying a rotation by a set $\boldsymbol{\rho} = (\alpha, \beta, \gamma)$ of Euler angles of the wavelet Ψ^j initially located at the north pole, before computing an inner product between the wavelet and the field I:

$$(I \star \Psi^j)(\boldsymbol{\rho}) = \int_\Omega I(\boldsymbol{\omega})[R_\rho \Psi^j(\boldsymbol{\omega})]^* \mathrm{d}\Omega, \tag{8}$$

where R_ρ is the rotation by Euler angles $\boldsymbol{\rho}$, and $*$ stands for complex conjugation. From $(I \star \Psi^j)(\boldsymbol{\rho})$ we can identify (β, α) with the spherical coordinates $\boldsymbol{\omega} = (\theta, \varphi)$ and γ to the orientation that is probed in the convolution. In this way, we obtain oriented wavelet coefficients, with shorthand notation

$$(I \star \Psi^{j,\gamma})(\boldsymbol{\omega}) \equiv (I \star \Psi^j)(\alpha = \varphi, \beta = \theta, \gamma). \tag{9}$$

While $(I \star \Psi^{j,\gamma})$ is a convenient notation, that also matches with previous work, we however emphasize that there is no $\Psi^{j,\gamma}$ oriented wavelet by itself.

In practice, the directional convolution can be computed accurately and efficiently in Wigner space, which is the Fourier space associated with three-dimensional rotations described by Euler angles. In this space, the directional convolution between a field I and a wavelet Ψ^j yields (McEwen et al. 2013, 2015b)

$$(I \star \Psi^j)^\ell_{mn} = \frac{8\pi^2}{2\ell + 1} I_{\ell m} \Psi^{j*}_{\ell n}, \tag{10}$$

where $I_{\ell m}$ and $\Psi^j_{\ell n}$ are the spherical harmonic coefficients of I and Ψ^j, respectively, and $(\Psi^j \star I)^\ell_{mn}$ are the Wigner coefficients of the convolved field, i.e. the Fourier representation of the directional wavelet coefficients defined over Euler angles $\boldsymbol{\rho} = (\alpha, \beta, \gamma)$. To return to the spatial domain, we compute an inverse Wigner transform, defined as

$$(I \star \Psi^{j,\gamma})(\boldsymbol{\omega}) \equiv (I \star \Psi^j)(\boldsymbol{\rho}) = \sum_{\ell=0}^{L} \frac{2\ell + 1}{8\pi^2} \sum_{m,n=-\ell}^{\ell} (I \star \Psi^j)^\ell_{mn} D^{\ell *}_{mn}(\boldsymbol{\rho}), \tag{11}$$

where $D^\ell_{mn}(\boldsymbol{\rho})$ are the Wigner-D matrices (Varshalovich et al. 1988). Fast (inverse) Wigner transform algorithms can then be leveraged for efficient computation (McEwen et al. 2015a; Price & McEwen 2024). By computing the wavelet transform through harmonic space as described, pixelisation artefacts are avoided. Although the wavelets satisfy an admissibility condition such that the field can be recovered exactly from its wavelet coefficients (McEwen et al. 2015b), we are only concerned with the forward wavelet transform in this work. In the following, we also group the scale and orientation under a single index $\lambda = (j, \gamma)$, writing $I \star \Psi^\lambda$ for the wavelet transform of the field I at a given oriented scale λ.

While computing convolutions through harmonic representations is highly accurate it involves moderate computational cost since generalised Fourier transforms on the sphere and space of rotations must be computed (albeit using fast algorithms). An alternative would be to compute the convolutions in pixel space as done in Delouis et al. (2022). However this can introduce pixelisation artefacts. A future avenue to consider is hybrid discrete-continuous approaches as shown to be highly computationally efficient while also avoiding discretization artefacts (Ocampo et al. 2023).

In order to optimize our numerical implementation, in particular the memory usage, all convolutions are computed in a multi-scale framework, where the map resolution is tuned to the scale at which the convolution is made. See Leistedt et al. (2013) for more details.

Fig. 2 illustrates orthographic projections of the directional spherical wavelets for two scales and three angles, viewed looking down from the North pole.

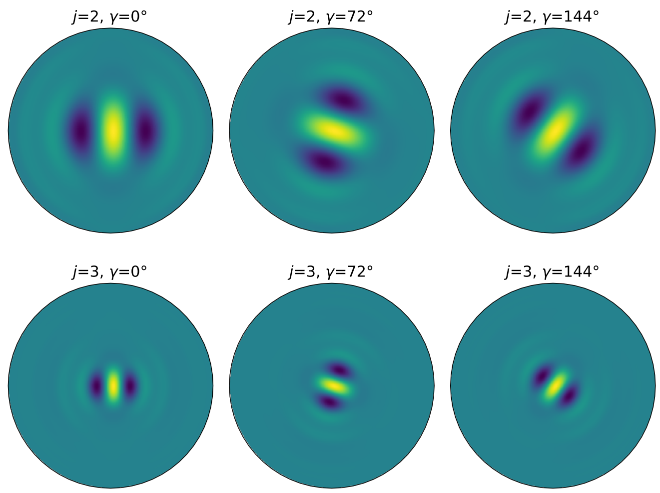

Fig. 2. Orthographic projections of the directional spherical wavelets for two scales and three angles, viewed looking down from the North pole. In this case, we set $N = 3$ so that the angle γ takes $2N - 1 = 5$ values between 0 and 360 deg but only three are shown here.

2.3. Scattering covariance coefficients

Scattering transforms cover several types of summary statistics, see for instance Allys et al. (2019, 2020); Cheng et al. (2024b). In this work, we consider the scattering covariances, or scattering spectra, previously introduced in Morel et al. (2023) and Cheng

[1] Although steerability could be exploited in future for further computational savings (McEwen et al. 2015b).

et al. (2024b). We chose the SC, because they only rely on convolutions, and not on translations as the Wavelet Phase Harmonics (Allys et al. 2020), which are difficult to univocally define on spherical maps.

SC statistics characterize the power and sparsity at given scales, as well as interaction between different scales. They are built from successive applications of wavelet transforms and modulus operators, followed by average and covariance computations (assuming that we work with homogeneous processes). For a detailed introduction to these statistics, the reader is invited to refer to Cheng et al. (2024b). We consider two coefficients at a single scale j_1 and a single angle γ_1, *i.e.* $\lambda_1 = (j_1, \gamma_1)$:

$$S_1^{\lambda_1} = \langle |I \star \Psi^{\lambda_1}| \rangle , \tag{12}$$

$$S_2^{\lambda_1} = \langle |I \star \Psi^{\lambda_1}|^2 \rangle , \tag{13}$$

and two coefficients that characterize the couplings between two and three oriented scales[2]:

$$S_3^{\lambda_1,\lambda_2} = \mathrm{Cov}\left[I \star \Psi^{\lambda_1}, |I \star \Psi^{\lambda_2}| \star \Psi^{\lambda_1} \right], \tag{14}$$

$$S_4^{\lambda_1,\lambda_2,\lambda_3} = \mathrm{Cov}\left[|I \star \Psi^{\lambda_3}| \star \Psi^{\lambda_1}, |I \star \Psi^{\lambda_2}| \star \Psi^{\lambda_1} \right], \tag{15}$$

where $\langle\cdot\rangle$ corresponds to the mean over the sphere, defined in Sec. 2.1, and where covariances are defined as $\mathrm{Cov}[XY] = \langle XY^* \rangle - \langle X \rangle \langle Y^* \rangle$ for two complex fields X and Y. Note that in our case, the wavelet transforms are always zero mean, since the wavelets are mean-free. Since taking the modulus of a field mainly displaces its frequency support toward lower frequency (Zhang & Mallat 2021; McEwen et al. 2022), it is sufficient to consider terms for which $j_1 \le j_2 \le j_3$. Moreover, we introduce the additional parameter δ_j, which corresponds to the maximum distance between pairs of scales whose interactions is characterized: *i.e.*, the calculation of S_3 and S_4 are restricted to pairs of scales such that $j_2 - j_1 \le \delta_j$ and $j_3 - j_1 \le \delta_j$.

The dimension of S_1 and S_2 coefficients is $J\Theta$ with J the number of scales and $\Theta = 2N - 1$ the number of orientations. Regarding S_3 and S_4, their dimensions are approximately $J^2\Theta^2$ and $J^3\Theta^3$ or $J\delta_j\Theta^2$ and $J\delta_j\Theta^3$ when considering a maximum distance between scales δ_j. The exact number of coefficients are given in Tab. 2.

The power spectrum of the field is mainly characterized by the $S_2^{\lambda_1}$ coefficients defined in Eq. 13. These terms correspond to the average of the power spectrum over the ℓ-wavelet bandpasses (Cheng et al. 2024b). However, they only constrain the power spectrum over a small number of bands and this is usually not precise enough for modeling purpose. Increasing the number of scales j that we probe can be done by decreasing the wavelet dilation parameter η. However, this leads to a large increase of the total number of SC coefficients. For this reason, we have considered additional $S_2^{\lambda_1}$ coefficients, built with a second filter set with $\eta' < \eta$ ($\eta = 2$ and $\eta' \simeq 1.58$) and a single orientation $N' = 1$ (isotropic filters). These coefficients are called $S_2^{\lambda_1'}$ and they allow us to constrain the power spectrum over thinner ℓ-bins.

For physics fields the power spectra can typically be modelled by a power law, at least over certain scales (Cheng et al. 2024b). This leads to SC coefficients which can vary over several orders of magnitude, since their amplitude is controlled by the $I \star \Psi^{\lambda}$ terms, which filters the initial I field over the j frequency band of the wavelet. This amplitude discrepancy can lead to ill-conditioned optimizations. Following previous works (see, for instance Cheng et al. 2024b), we avoid this issue by normalizing the SC statistics from the S_2 coefficients of a reference field that we note $S_{2,\mathrm{ref}}$. We thus define

$$\bar{S}_1^{\lambda_1} = \frac{S_1^{\lambda_1}}{\sqrt{S_{2,\mathrm{ref}}^{\lambda_1}}}, \quad \bar{S}_2^{\lambda_1} = \frac{S_2^{\lambda_1}}{S_{2,\mathrm{ref}}^{\lambda_1}}, \quad \bar{S}_2^{\lambda_1'} = \frac{S_2^{\lambda_1'}}{S_{2,\mathrm{ref}}^{\lambda_1'}}, \tag{16}$$

$$\bar{S}_3^{\lambda_1,\lambda_2} = \frac{S_3^{\lambda_1,\lambda_2}}{\sqrt{S_{2,\mathrm{ref}}^{\lambda_1} S_{2,\mathrm{ref}}^{\lambda_2}}}, \quad \bar{S}_4^{\lambda_1,\lambda_2,\lambda_3} = \frac{S_4^{\lambda_1,\lambda_2,\lambda_3}}{\sqrt{S_{2,\mathrm{ref}}^{\lambda_2} S_{2,\mathrm{ref}}^{\lambda_3}}} . \tag{17}$$

When constructing generative models below, we will take the target field as the reference field, which will allows us to deal with coefficients whose values will be at most of order unity.

3. Generative modeling

In this section we describe how to build generative models from the SC statistics of a given field. We also give some details on the numerical implementation of the generative models and the associated computational cost.

3.1. Maximum entropy generative model

We build generative models under SC constraints. These are maximum entropy microcanonical models, which are approximately sampled by gradient descent. We refer the readers to Bruna & Mallat (2019) for more details.

Such models are constructed from statistics Φ estimated from a target field x_t; in this paper the target field is a single full-sky map. The associated microcanonical set Ω_ε of width ε is

$$\Omega_\varepsilon = \left\{ x : ||\Phi(x) - \Phi(x_t)||^2 < \varepsilon \right\}, \tag{18}$$

where $||.||$ is the Euclidean norm. The microcanonical maximum entropy model is the model of maximal entropy defined over Ω_ε, which has a uniform distribution over this set.

In this paper, we approximate this sampling with a gradient descent approach, which consists in transporting a higher entropy white Gaussian distribution into a distribution supported in Ω_ε. In practice, each new sample is obtained by first drawing a white noise realization, and then performing a gradient descent in pixel space, or in harmonic space, using a loss function

$$\mathcal{L}(x) = ||\Phi(x) - \Phi(x_t)||^2 . \tag{19}$$

The typical width ε of the microcanonical ensemble is then fixed by the number of iterations used in the gradient descent. The numerical details of this implementation are given in Sec. 3.2.

In our case, the summary statistics $\Phi(x)$ that we consider are the mean over pixels $\langle x \rangle$, its variance $\mathrm{Var}(x)$ and the normalized SC statistics, defined in Sec. 2.3. Thus, we have

$$\Phi(x) = \left\{ \langle x \rangle, \mathrm{Var}(x), \bar{S}_1^{\lambda_1}, \bar{S}_2^{\lambda_1}, \bar{S}_2^{\lambda_1'}, \bar{S}_3^{\lambda_1,\lambda_2}, \bar{S}_4^{\lambda_1,\lambda_2,\lambda_3} \right\}. \tag{20}$$

We note that the target statistics $\Phi(x_t)$ are evaluated from a single full-sky image, and that the SC generative models are then

[2] Note than in Cheng et al. (2024b), the $S_3^{\lambda_1,\lambda_2}$ are defined as

$$S_3^{\lambda_1,\lambda_2} = \mathrm{Cov}\left[I \star \Psi^{\lambda_1}, |I \star \Psi^{\lambda_2}| \right].$$

However, as only the ℓ harmonics appearing in both sides of the covariance have a non-vanishing contribution, only harmonics captured by λ_1 of the $|W^{\lambda_2} I|$ term play a non-negligible role, and both formulations are closely related.

built from this single set of constraints. In this respect, our approach differs from machine learning-based approaches, which generally require training on a large, and potentially very expensive, dataset.

3.2. Details on the numerical implementation

In this work we consider scattering covariances defined on the sphere, constructed using spherical wavelet transforms which in turn depend on efficient spherical harmonic and Wigner transforms (see Sec. 2.2). As outlined in Sec. 3.1 new images are drawn from a SC microcanonical model by minimising the loss defined in Eq. 19. A plethora of algorithms have been developed to solve such optimisation problems, however they typically require gradient information, which in turn requires that each component of the loss be differentiable. Consequently, we require the spherical scattering covariance and thus the spherical wavelet, spherical harmonic and Wigner transforms to all be differentiable. Recently, open-source `JAX` software that is differentiable and GPU accelerated has been developed for all of these transforms, including `s2fft`[3] for spherical harmonic and Wigner transforms (Price & McEwen 2024) and `s2wav`[4] for spherical wavelet transforms (Price et al. 2024). As part of the current work, we have developed a new open-source software implementing the spherical scattering covariances called `s2scat`[5], which builds on top of `s2fft` and `s2wav`.

For a given target field we compute a generated field by minimising the loss defined in Eq. 19 through a gradient descent in harmonic space with different initial conditions. These initial conditions are Gaussian white noises sampled in spherical harmonic domain, *i.e.*, whose all $I_{\ell m}$ real and imaginary parts are drawn from the same Gaussian distribution such that the total variance of the target field is reproduced. In this way, the starting angular power spectrum, as defined in Eq. 21, is flat.

To avoid a repeated spherical harmonic transform as the first step at each iteration in the computations, we chose to perform the gradient descent in the spherical harmonic domain rather than in pixel space. The variables we iterate on during the loss minimization thus are the $I_{\ell m}$ coefficients. Because the maps are real and thanks to relation (3), we can iterate on the $I_{\ell m}$ with $m \geq 0$ only. The loss minimization is done through a gradient descent with the L-BFGS algorithm described in Byrd et al. (1995), using the `JAX` auto-differentiable framework (Bradbury et al. 2018) and the `jaxopt` package (Blondel et al. 2022). We stop the optimization after ~ 400 iterations, which in our experiment is the typical time for the loss function to decrease by about four orders of magnitude and to reach a plateau at values around 0.1 (meaning, since the loss is not normalized by the number of coefficients, that coefficients are on average constrained at subpercent accuracy).

3.3. Computational cost

As outlined in Sec. 2, computation of the scattering covariance statistics requires repeated spherical convolutions with subsequent non-linear activation functions, in this case modulus operators. Although directional spherical convolutions may naïvely be computed in pixel space with complexity $O(L^5)$ (McEwen et al. 2007), they are more efficiently evaluated in harmonic space with complexity $O(NL^3)$ (McEwen et al. 2007, 2013, 2015b). Furthermore, excellent accuracy can be achieved by computing convolutions in harmonic space since pixelisation artefacts are avoided.

Precompute Mode			
Bandlimit	Forward	Gradient	JIT Compilation
256	15 ms	30 ms	20 s
512	100 ms	200 ms	25 s
Recursive Mode			
Bandlimit	Forward	Gradient	JIT Compilation
256	120 ms	300 ms	90 s
1024	5 s	10 s	3 m
2048	20 s	50 s	6 m

Table 1. Computational benchmarking results of the scattering covariance transform provided by `s2scat`. These results were recovered on a single NVIDIA A100 40GB GPU, although it is possible to run across multiple GPUs. In our analysis we generate spherical images through 400 iterations to be conservative. In practice however, we find that $\sim$ 100 iterations is typically sufficient, in which case an image at $L = 256$ may be generated in $\sim$ 4s. Furthermore, batched generation can dramatically decrease per sample compute time. For example, 20 images at $L = 256$ can be generated in $\sim$ 12s, corresponding to $\sim$ 0.5s per sample.

We must repeatedly map to and from spherical harmonic space within our generative model using `s2fft` (Price et al. 2023). Two operating modes are provided by `s2scat`: one computes and caches the reduced Wigner d-functions, which are then used at runtime (precompute mode); and the other computes these functions on-the-fly through recursive algorithms (on-the-fly mode). Conceptually, the precompute mode is fast but requires $O(L^3)$ memory, whereas the on-the-fly mode is slower but requires at most $O(L^2)$ memory. When running on GPUs for harmonic bandlimits $L \leq 512$ we recommend one adopts the precompute mode, deferring to the on-the-fly algorithms at higher resolutions. Although with GPU memory increasing rapidly with hardware developments, it is likely that the precompute mode can be ran at higher bandlimits on the latest and upcoming GPUs.

High-level benchmarking results are presented in Tab. 1. In each case we consider an azimuthal bandlimit of $N = 3$, which corresponds to 5 directions on the sphere, and the full set of anisotropic scattering covariances. Our benchmarking was performed on a single NVIDIA A100 GPU with 40GB of on board memory, though in practice `s2scat` may be distributed across a large number of GPUs. For completeness, we record the time for both a forward and gradient evaluation, in addition to the time required for Just-In-Time (JIT) compilation. One may also utilise the `jax.vmap` API, which allows one to batch calls to the maximum entropy model presented in Sec. 3, resulting in more optimal GPU utilisation. For example, suppose we sample from our microcanonical model with 100 iterations of a first order optimiser (*e.g.* ADAM; Kingma & Ba 2014). Generating a single new image at $L = 256$ takes $\sim$ 4s, whereas a batched call to generate 20 such images takes $\sim$ 12s which is $\sim$ 0.5s per new sample. Furthermore, the `jax.pmap` API allows one to batch calls across GPU devices, therefore accelerating generation linearly with the number of available GPUs.

4. Validation of the generative models

In this section we construct SC generative models from four astrophysical fields showing different types of structures. We then compare the generated fields with the target one. We first do a visual comparison of the maps to asses the quality of the spatial

[3] https://github.com/astro-informatics/s2fft
[4] https://github.com/astro-informatics/s2wav
[5] https://github.com/astro-informatics/s2scat

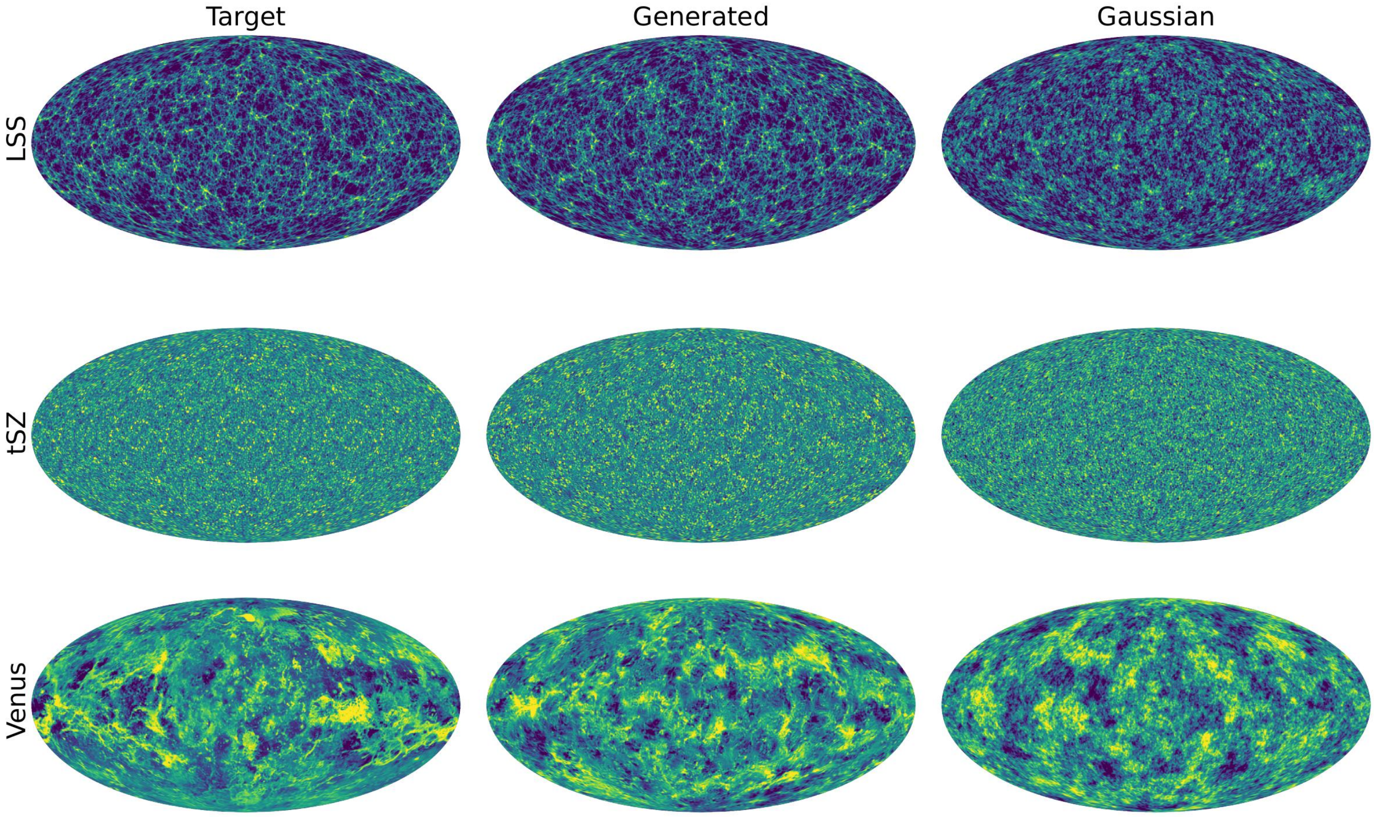

Fig. 3. Maps for LSS, tSZ and Venus fields (from top to bottom). The left column is the original target field. The central column shows one sample of the generated maps. The right column shows a Gaussian field with the same power spectrum as the target. For LSS and tSZ, we plot the logarithm of the fields, in order to better see the texture details. Color bars are identical within each field.

texture reproduction. We then compute various statistics in order to quantitatively evaluate our generative model. For each type of data, we draw 50 samples of the microcanonical ensemble and compute the mean and the standard deviation over those 50 realisations. For the LSS and the tSZ fields, while the optimization is performed on the logarithm of the maps, we compare the statistics on the raw image, after taking the inverse transform.

Before any comparison, all maps are filtered in harmonic space, only keeping the $I_{\ell m}$ coefficients such that $\ell_{\min} \leq \ell \leq L$ with $\ell_{\min} = \eta^{J_{\min}}$ the central frequency of the wavelet associated to $J_{\min}$. In this way, we consider only the scales that are constrained during the optimization. Note however that if necessary, it is possible to constrain in a similar scheme scales up to $l_{\min} = 0$.

We also propose a comparison with samples from a Gaussian model built from the power spectrum of the target field. We have produced 50 Gaussian realisations using the `synfast` method from the `Healpix` package (Górski et al. 2005) which allows one to construct such a Gaussian model from the angular power spectrum of a target field. This allows us to quantify the contribution of our models built from SC statistics compared to purely Gaussian statistics.

In Appendix A, we directly compare the SC statistics of the target and the generated fields. This is an additional check for the quality of the generative model. Indeed we expect them to match as they are part of the coefficients constrained during the optimization.

4.1. Description of the set of maps

We first present the astrophysical and cosmological fields from which we construct SC generative models. They are expected to have homogeneous statistical properties on the sphere. This property is essential since this is assumed when computing statistics through spatial averages. This requirement, as well as a possible way to avoid this constraint, is discussed in Sec. 4.4. The four different fields, denoted as follow, are:

- **LSS**: a Large Scale Structure simulation of weak lensing, from the CosmoGrid data set (Kacprzak et al. 2023; Fluri et al. 2022);
- **tSZ**: a thermal Sunyaev-Zeldovitch effect simulated map from Simons observatory Galactic foreground simulations (Ade et al. 2019), which were produced using the Sehgal et al. (2010) model with modifications to better match the recent measurements;
- **Venus**: a map of the Venus planet from Science On a Sphere database[6];
- **CMB**: a CMB temperature map produced using the Python Sky Model software (`PySM`) from Thorne et al. (2017).

We refer the readers to the references given above for more details on these fields. As a Gaussian field, the CMB map is a good null-test for our method, as presented in Appendix B. The other fields all originate from non-linear physical processes, and thus have highly non-Gaussian structures, as can be seen in the left column of Fig. 3. The diversity of these fields illustrates the generality and the versatility of our method, which could be used for various physical data sets.

[6] `https://sos.noaa.gov/sos/`

	J	$J_{\min}$	N	δ_j	Trans.	Nb. of terms
LSS	7	2	3	5	log.	8283 (35, 450, 7750)
tSZ	6	3	3	5	log.	6173 (30, 350, 5750)
Venus	7	2	3	5	lin.	8283 (35, 450, 7750)
CMB	8	1	3	5	lin.	10393 (40, 550, 9750)

Table 2. For each field, we give the number J of scales that we probe, the value of $J_{\min}$, the value of N which corresponds to the angular resolution of the wavelets and the δ_j parameter which corresponds to the maximum distance between pairs of scales whose interactions is characterized. We also indicate if we perform a logarithmic (log.) or a linear (lin.) transformation on the target map, as discussed in Sec. 4.1. The last column gives the total number of terms that composes our summary statistics $\Phi(x)$, with, between parenthesis, the detailed count for S_1 (equal to S_2), S_3 and S_4.

While all these simulated maps are available in Healpix format, the computation of the SC is done directly from the harmonic space. The conversion to this space is done using $L - 1 = \ell_{\max} = 2\texttt{nside}$, and thus acts as a low-pass filter operation. This implies that spatial frequencies at $\ell \geq L$ are filtered out in this operation. Note that calculating the SC and performing the optimization directly in spherical harmonic space means that there is no particular constraint on the sampling on the target map, even if the internal SC calculation steps are based here on MW or alternative sampling schemes to improve accuracy.

During the optimization process, all maps are normalized such that their mean is zero and their standard deviation is one. In addition, the LSS and tSZ fields are highly non-Gaussian, making them difficult to model directly even with SC statistics. This is why we instead chose to model the logarithm of these maps[7]. This logarithmic transform brings the distribution closer to a Gaussian one, and reduces in particular the weight of the high amplitude tail of the probability density function, allowing for better SC generative models. At the end of the optimization, we however take the inverse transform for these maps, and we assess the quality of the generative model on the raw images.

The generative models have been run on MW maps with a resolution of $L = 256$ which corresponds to $L(2L - 1) = 130816$ pixels. These real signals have $L^2 = 65536$ complex harmonic coefficients. For the directional wavelets used to build the SC, we have considered a dyadic scaling $\eta = 2$ and $N = 3$. This leads to $J_{\max} = 8$ and $2N - 1 = 5$ orientations. As shown in Tab. 2, the value of $J_{\min}$ was tuned to each field in order to take into account the largest spatial scales that compose the maps. The number J of scales that we probe is also given in the Table. The maximum distance δ_j between scales was fixed to five, which allows us to divide the total number of SC coefficients by approximately two without degrading the quality of the generative model. Concerning the additional $S_2^{\lambda'}$ coefficients, we chose an axisymmetric filter set with $N' = 1$ and a scaling given by $\eta' \simeq 1.58$. The exact total number of terms that composes the summary statistics $\Phi(x)$ is given in Tab. 2.

4.2. Visual validation

As a first test we can visually compare the target field and the generated ones, as presented in Fig. 3 for the LSS, tSZ and Venus fields. They appear to be visually very similar to the original maps, which clearly shows that the SC statistics capture an important part of the non-Gaussian texture of the field. On the contrary, the structures are not reproduced in the Gaussian realisations shown on the right column. For LSS and tSZ, we plot the logarithm of the fields which allows to better see the textures. In addition, in Fig. 4, we show a zoom on a smaller region to better visualize the details in the spatial structures. In Appendix C, we also show four realisations of the fields starting from different initial conditions. This shows the ability of our generative models to sample independent realizations while capturing the overall texture of the fields.

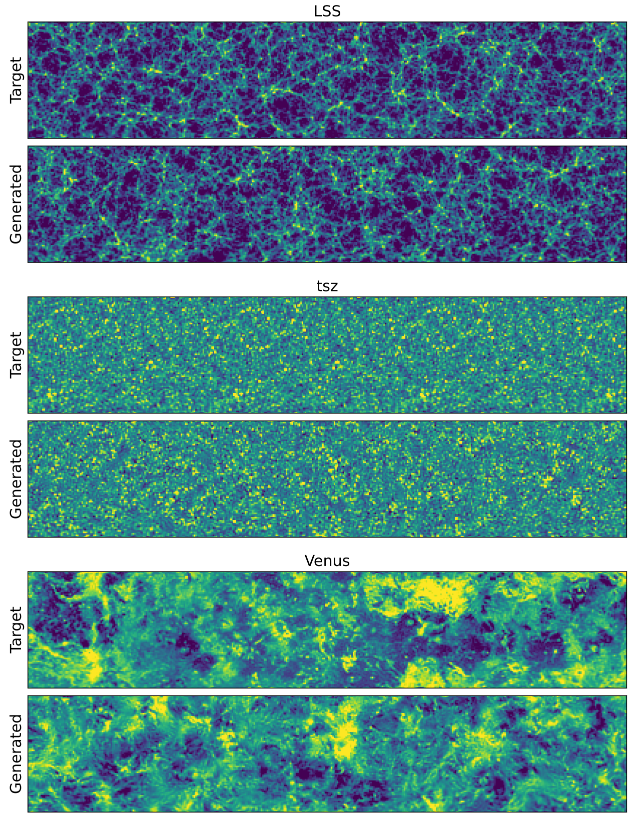

Fig. 4. Zoom on a region for the LSS (top), tSZ (middle) and Venus (bottom) fields in order to see details of the texture. Similarly to the previous figure, for LSS and tSZ, we plot the logarithm of the fields, in order to better see the texture details. Color bars are identical within each field.

4.3. Statistical validation with standard summary statistics

Following a similar approach to Cheng et al. (2024b) and Price et al. (2023), we compare summary statistics between the target field and the generated field. As previously, we show the mean and the standard deviation computed over the 50 realisations. The summary statistics we chose to compare are:

- the Probability Density Function (PDF) of the map;
- the angular power spectrum;
- the three Minkowski functionals.

The PDF of the maps and the three Minkowski functionals are performed on Healpix maps. To do that, the output $I_{\ell m}$ from the loss minimization are projected to Healpix map by an inverse spherical harmonic transform at the end of the generative process.

[7] For the LSS field, the exact logarithmic transform applied on I is $\log(I + \varepsilon)$ where $\varepsilon = 10^{-3}$ is a regularization to deal with non-positive fields.

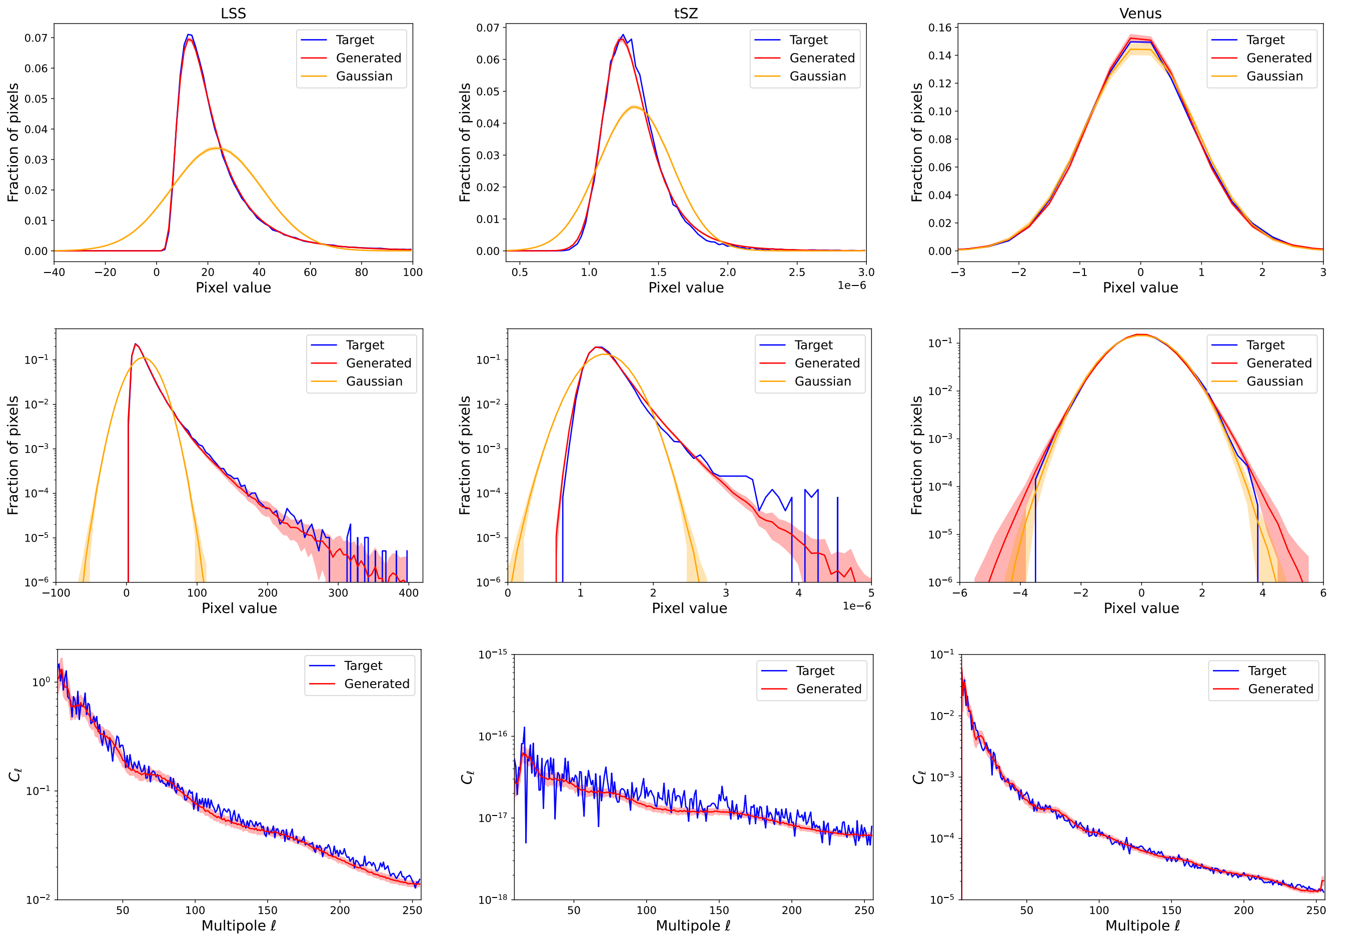

Fig. 5. PDF and angular power spectra for the LSS, tSZ and Venus fields, from left to right. The first row shows the PDF with a linear y-axis scaling while the second row shows the same PDF with a logarithmic y-axis. The third row shows the angular power spectra. The target is shown in blue, the generated fields in red, and the Gaussian realisations in yellow. We plot the mean (solid line) and the standard deviation (shadow envelope) over 50 realisations.

4.3.1. Probability density function

The PDF for the LSS, tSZ and Venus fields are shown in Fig. 5, computed on the Healpix maps. On the first row, we show the PDF with a linear y-axis scaling, while on the second row, we show them with a logarithmic y-axis scaling in order to better exhibit the tails of the distributions on several orders of magnitude. The target fields are shown in blue and the generated ones in red. In yellow, we also show the comparison with the Gaussian realizations. By definition, the PDF of these realizations presents a Gaussian profile in linear scale and a parabolic profile in logarithmic scale.

While the Venus field has a PDF which only slightly differs from the Gaussian case, the two other fields clearly have non-Gaussian features with large tails. The comparison of the target and generated PDFs with the Gaussian PDF also allows us to better see their non symmetric shape, which is characteristic of non-Gaussian features. As we can see, the PDFs for SC models are well reproduced on at least three orders of magnitude. The results obtained for the LSS fields, which are very good up to five orders of magnitude, are especially striking. On the other hand, results for the tSZ field begins begins to push the expressive limit of our generative models. For the Venus maps, we identify the abrupt jump in the histogram as a flaw in the data used, which does not particularly illustrate a limitation of our maximum entropy SC model.

4.3.2. Angular power spectrum

We calculate the power spectrum in the usual way as

$$C_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{m=\ell} |I_{\ell m}|^2 \,, \tag{21}$$

where the normalization factor $1/(2\ell + 1)$ yields a flat power spectrum in case of white Gaussian noise.

As discussed in Sec. 2.3, the power spectrum is constrained during the optimization through the S_2 and the additional S_2' coefficients. Note however that these coefficients do not constrain the full power spectrum, each term only constraining a weighted power spectrum over the frequency support of the associated wavelet.

The third row of Fig. 5 shows the results for the generation of LSS, tSZ and Venus fields, from left to right. Power spectra are well reproduced over all scales, even when they vary over up to four orders of magnitude. However, small oscillations around the target can be seen in the generated power spectra. These are residual features related to the frequency bands of the wavelets, which illustrate the trade-off between the quality of reproduction we want to achieve and the number of filters we use; i.e. the computational efficiency of our generative model.

Note that in this paper, we include 11 of these additional S_2' coefficients. This number, and the precise shape of the wavelets

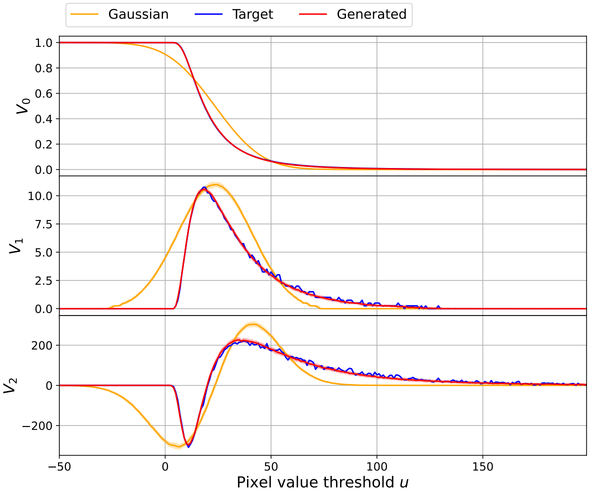

Fig. 6. The three Minkowski functionals V_0, V_1 and V_2 for the LSS field. Blue is the target, red the generated field and yellow the Gaussian fields. For the generated field and the Gaussian realisations, we plot the mean (solid) and the standard deviation (shadow envelope) computed over 50 realisations.

used, could however be tuned to better reproduce the power spectrum of the target. However, care must be taken not to overconstrain these terms, as all the samples generated would then have a power spectrum very close to that of the target, which does not necessarily correspond to a good generative model. The introduction of S_2' terms is an improvement with respect to previous work, since it allows us to have better power spectrum constraints without significantly increasing the overall number of SC statistics.

4.3.3. Minkowski functionals

Finally, we compute the Minkowski functionals on the Healpix maps. These standard non-Gaussian statistics characterize the topology of the level sets of the field. In two dimensions, there are three Minkowski functionals $V_0(u)$, $V_1(u)$ and $V_2(u)$ which depend on a pixel value threshold u. We computed them with `Pynkowski` software (Carones et al. 2024). We refer the reader to this publication for the complete definition of those statistics. The result is shown in Fig. 6 only for the LSS field while the others are shown in Appendix D. In yellow, as a comparison, we show the case of the Gaussian realisations. For the generated field and the Gaussian realisations, we plot the mean (solid) and the standard deviation (shadow envelope) computed over 50 realisations. Thus, the SC models encompass very well these non-Gaussian statistical features.

4.4. Limitations and discussions

This work is a first implementation of generative models from state-of-the-art spherical SC. As a first proof of concept, we constructed and validated these models on various cosmological fields, most of them simulated. However, building such models from real data, or using these tools to perform statistical component separation, may require some additional work to deal with their own specificity. In this section, we comment some of these limitations, and how to overcome them.

A first limitation of our current work is the fact that our models assume the statistical homogeneity of the fields studied. However, the ability of dealing with non-homogeneous physical processes is usually required when modeling astrophysical fields, such as the Galactic emissions, which properties typically vary strongly with latitude on the sky. An efficient way to deal with this issue is to rely on different masks in pixel space, for which statistical properties can be constrained independently (Delouis et al. 2022). However, this requires a trade-off between the size and number of masks: using a larger number of masks gives a better description of large-scale variations in statistical properties, but increases the variance of SC statistics estimates on each mask due to the smaller number of pixels used, as well as the total number of SC coefficients and the computational and memory cost.

A second limitation is the map resolution we can achieve. For now, the generation of a new map at $L = 256$ and $N = 3$ takes $\leq$ 1s on a single GPU. This is thanks in part to a large number of precomputed matrices necessary for the Wigner transform which are, stored in memory (several Gigabytes). Specifically this memory if cubic with L, which is prohibitive at high L. When increasing the resolution beyond $L \sim 1024$, we usually reach the GPU memory limit and the coefficients have to be computed on-the-fly, though this of course depends on, available GPU specifications. Critically, the on-the-fly approach dramatically reduces memory requirements so that generations at high L are at least feasible but at the cost of a significantly increased computation time. In future we will explore further optimizations for high L. A key avenue we will explore is the introduction of hybrid wavelet convolutions which operate efficiently in pixel and harmonic space at high and low resolutions respectively (see *e.g.* Delouis et al. 2022; Ocampo et al. 2023).

5. Conclusions

The main result of this paper is the extension of state-of-the-art scattering transforms to spherical fields. We have worked with the last generation of scattering transform statistics, named scattering covariances (Cheng et al. 2024b), which were previously introduced for one dimension and two dimensions planar fields. They have the advantage of relying only on successive wavelet transforms and modulus, as well as on covariances, and do not require any translations. We have also used state-of-the-art directional wavelets on the sphere, computed in spherical harmonic space (McEwen et al. 2015b). The numerical implementation of this work `s2scat`, is open-source and publicly available. Furthermore, it is fully autodifferentiable, using the JAX Python framework (Bradbury et al. 2018) and building on the `s2fft`/`s2wav` packages (Price & McEwen 2024; Price et al. 2024).

These developments allow us to build generative models of full sky spherical fields without the need for large training datasets. In fact, our method holds in the limit of a single data realisation. The performance of those generative models were validated quantitatively on different fields: a LSS weak lensing field, maps of tSZ effect and of the CMB, as well as a map of Venus surface, for which they performed extremely well. The diversity in terms of structures between the maps shows the impressive ability of SC to comprehensively characterize very different non-Gaussian textures.

This work introduces a new powerful innovative approach for spherical data, and it opens interesting perspectives for astrophysical applications. In particular, we plan to use it for the study and the modeling of CMB astrophysical foregrounds. The first goal will be to have a tool to produce multiple realisations of the different astrophysical components, for example using the AGORA simulations (Omori 2024). Then, scattering transforms could play a role in component separation, relying both on recently developed scattering transform-based statistical component separations approaches, as well as investigating how classical component separation methods could benefit from scattering transforms, using the non-Gaussianities as an additional lever arm to disentangle different components.

Finally, we also point out that SC statistics provide highly informative sets of statistics, which could be very useful for tasks such as parameter inference, for instance Simulation Based Inference (SBI; Cranmer et al. 2020), from large cosmological surveys (see, for instance Régaldo-Saint Blancard et al. 2024; Gatti et al. 2024; Cheng et al. 2024a). This could be all the more useful as the compression factor they enable, compared with a direct description in pixel space, becomes extremely large at high resolution, due to logarithmic scale binning. This property could be further enhanced by using compression schemes like the one presented in Cheng et al. (2024b), which can make SC a very informative and versatile compressed set of statistics.

Acknowledgements. We thank Sixin Zhang for helpful discussions about scattering transforms. Most of the simulations were produced on the MesoPSL calculation center, for which we thank the administrators. MAP and JDM are supported in part by EPSRC (grant number EP/W007673/1) and STFC (grant number ST/W001136/1).

Appendix A: SC statistics

In Fig. A.1 we show the normalized SC coefficients $\bar{S}_1$, $\bar{S}_2$, $\bar{S}_3$ of the LSS field. We plot the coefficients of the target field in blue, the generated ones in red, as well as the Gaussian realisations in yellow. Coefficients are plot following the lexicographic order. We chose not to show $\bar{S}_4$ for readability because of the large number of coefficients.

Regarding the Gaussian realisations, shown in yellow, we expect the $\bar{S}_3$ coefficients to be equal to zero up to the correlations induced by the overlapping between wavelet bands. As we can see, for $\bar{S}_3$, the mean is centered on zero.

By construction, $\bar{S}_2$ for the target field is equal to one. This is because we have considered the S_2 coefficients of the target as the reference to normalize all the coefficients, as described in Sec. 2.3. In this way, all the normalized coefficients are of the order of the unit. As we can see, SC statistics are well constrained by the optimization, the generated coefficients in red well overlap the target coefficients in blue. $\bar{S}_3$ coefficients strongly differ from the Gaussian field, showing clear non-Gaussian signatures.

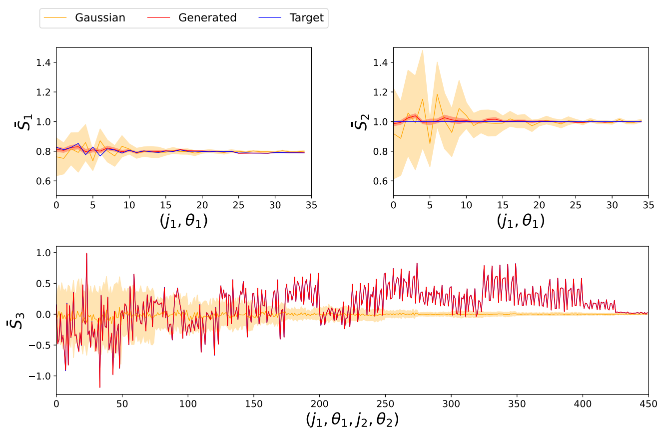

Fig. A.1. The normalized SC coefficients $\bar{S}_1$, $\bar{S}_2$, $\bar{S}_3$ for the logarithm of the LSS field. We show the coefficients from the target field (blue), the generated fields (red) and equivalent Gaussian realisations (yellow). The mean and the standard deviation over the 50 realisations are shown as a solid line with a shadow envelope.

Appendix B: CMB map as a null-test

It is important to check that the generative model behaves as we expect for a Gaussian field. This is an important validation for a maximum entropy generative model. This is why we tested it on the CMB map. The result is shown in Fig. B.1. The upper parts shows the target map, a generated field and a Gaussian realisation. As expected, the three maps look similar. We also plot the PDF and the power spectrum which match very well.

Appendix C: Multiple realisations

Fig. C.1 shows multiple realisations obtained from the generative model, changing the initial Gaussian random noise. For each fields, we show four maps out of the 50 that we computed. This is to illustrate the visual similarity between the realisations.

Appendix D: Minkowski functionals for the three others fields

Fig. D.1 shows the Minkowski functionals for the tSZ, Venus and CMB maps.

References

Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, J. Cosmology Astropart. Phys., 2019, 056
Allys, E., Levrier, F., Zhang, S., et al. 2019, A&A, 629, A115
Allys, E., Marchand, T., Cardoso, J. F., et al. 2020, Phys. Rev. D, 102, 103506
Auclair, C., Allys, E., Boulanger, F., et al. 2024, A&A, 681, A1
Blondel, M., Berthet, Q., Cuturi, M., et al. 2022, in Advances in Neural Information Processing Systems, ed. S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh, Vol. 35 (Curran Associates, Inc.), 5230–5242
Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: composable transformations of Python+NumPy programs
Bruna, J. & Mallat, S. 2013, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35, 1872
Bruna, J. & Mallat, S. 2019, Mathematical Statistics and Learning, 1, 257
Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, SIAM Journal on Scientific Computing, 16, 1190
Carones, A., CarrónDuque, J., Marinucci, D., Migliaccio, M., & Vittorio, N. 2024, MNRAS, 527, 756
Cheng, S., Marques, G. A., Grandón, D., et al. 2024a, arXiv preprint arXiv:2404.16085
Cheng, S. & Ménard, B. 2021, MNRAS, 507, 1012
Cheng, S., Morel, R., Allys, E., Ménard, B., & Mallat, S. 2024b, PNAS Nexus, 3, 103
Cheng, S., Ting, Y.-S., Ménard, B., & Bruna, J. 2020, MNRAS, 499, 5902
Cranmer, K., Brehmer, J., & Louppe, G. 2020, Proceedings of the National Academy of Sciences, 117, 30055
Delouis, J. M., Allys, E., Gauvrit, E., & Boulanger, F. 2022, A&A, 668, A122
Fluri, J., Kacprzak, T., Lucchi, A., et al. 2022, Phys. Rev. D, 105, 083518
Gatti, M., Jeffrey, N., Whiteway, L., et al. 2024, Physical Review D, 109, 063534
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Greig, B., Ting, Y.-S., & Kaurov, A. A. 2022, MNRAS, 513, 1719
Hothi, I., Allys, E., Semelin, B., & Boulanger, F. 2024, Astron. Astrophys., 686, A212
Jeffrey, N., Boulanger, F., Wandelt, B. D., et al. 2022, MNRAS, 510, L1
Kacprzak, T., Fluri, J., Schneider, A., Refregier, A., & Stadel, J. 2023, Journal of Cosmology and Astroparticle Physics, 2023, 050
Kingma, D. P. & Ba, J. 2014, arXiv preprint arXiv:1412.6980
Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv e-prints, arXiv:1110.3193
Lei, M. & Clark, S. E. 2023, ApJ, 947, 74
Leistedt, B., McEwen, J. D., Vandergheynst, P., & Wiaux, Y. 2013, Astron. & Astrophys., 558, 1
LiteBIRD Collaboration, Allys, E., Arnold, K., et al. 2023, Progress of Theoretical and Experimental Physics, 2023, 042F01
LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv e-prints, arXiv:0912.0201
Mallat, S. 2012, Communications on Pure and Applied Mathematics, 65, 1331
McEwen, J. D., Büttner, M., Leistedt, B., Peiris, H. V., & Wiaux, Y. 2015a, IEEE Sig. Proc. Let., 22, 2425
McEwen, J. D., Durastanti, C., & Wiaux, Y. 2018, Applied Comput. Harm. Anal., 44, 59
McEwen, J. D., Hobson, M. P., Mortlock, D. J., & Lasenby, A. N. 2007, IEEE Trans. Sig. Proc., 55, 520
McEwen, J. D., Leistedt, B., Büttner, M., Peiris, H. V., & Wiaux, Y. 2015b [arXiv:1509.06749]
McEwen, J. D., Vandergheynst, P., & Wiaux, Y. 2013, in Wavelets and Sparsity XV, SPIE international symposium on optics and photonics, invited contribution, Vol. 8858
McEwen, J. D., Wallis, C. G. R., & Mavor-Parker, A. N. 2022, in International Conference on Learning Representations (ICLR)
McEwen, J. D. & Wiaux, Y. 2011, IEEE Trans. Sig. Proc., 59, 5876
Morel, R., Rochette, G., Leonarduzzi, R., Bouchaud, J.-P., & Mallat, S. 2023, Scale Dependencies and Self-Similar Models with Wavelet Scattering Spectra
Ocampo, J., Price, M. A., & McEwen, J. 2023, in The Eleventh International Conference on Learning Representations (ICLR)
Omori, Y. 2024, Monthly Notices of the Royal Astronomical Society, 530, 5030
Price, M. A., Mars, M., Docherty, M. M., et al. 2023, The Open Journal of Astrophysics, 6, 35
Price, M. A. & McEwen, J. D. 2024, Journal of Computational Physics, 510, 113109
Price, M. A., Polanska, A., Whitney, J., & McEwen, J. D. 2024, arXiv e-prints, arXiv:2402.01282

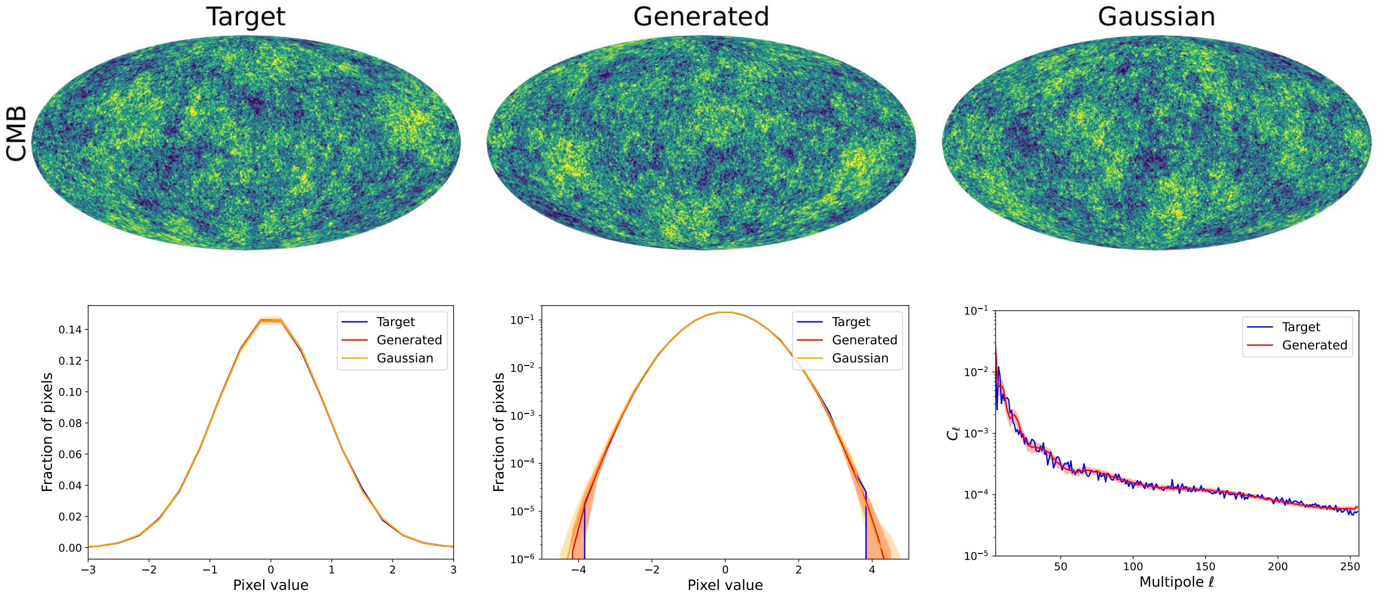

Fig. B.1. Result of the generative model for a CMB map. The upper part shows, from left to right: the target map, the generated field and a Gaussian realisation. Second row shows the PDF (linear and logarithmic scales) and the angular power spectrum.

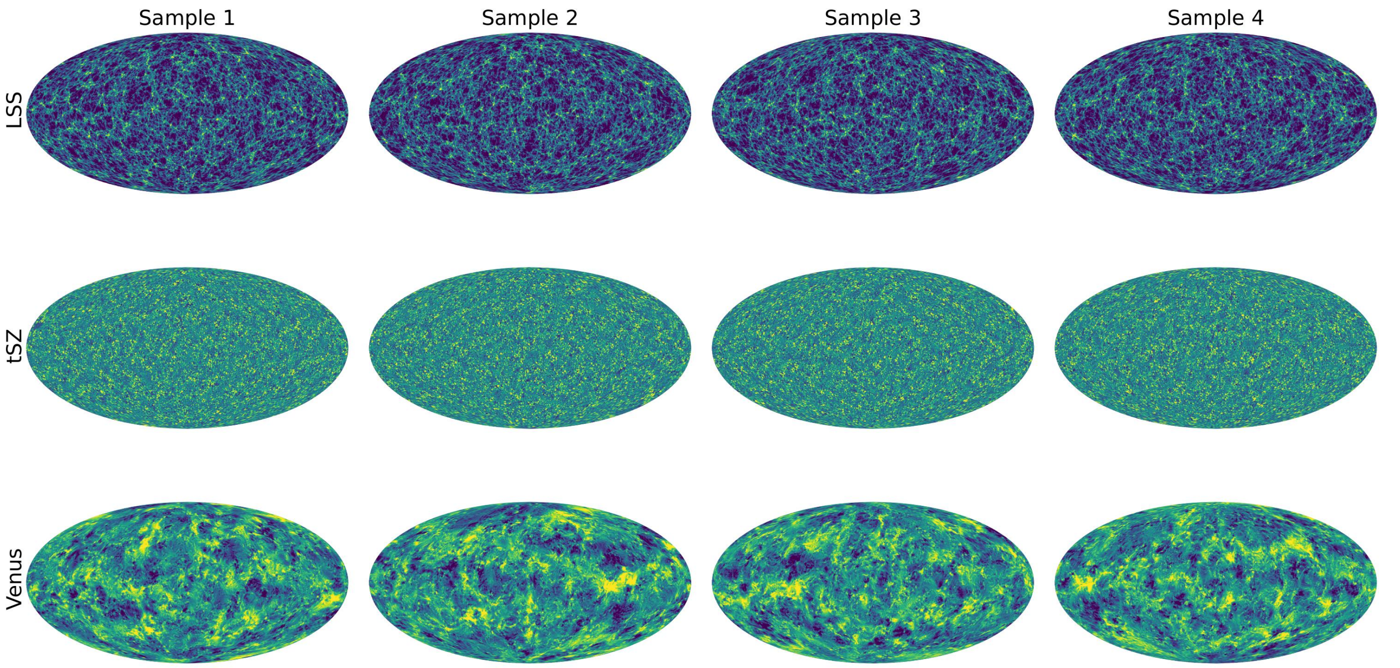

Fig. C.1. Four generative models of the LSS, tSZ and Venus fields (from top to bottom), obtained by changing the initial Gaussian random noise. In total we ran 50 realisations for each field. Color scales are identical within each field.

Régaldo-Saint Blancard, B., Allys, E., Auclair, C., et al. 2023, The Astrophysical Journal, 943, 9
Regaldo-Saint Blancard, B., Allys, E., Boulanger, F., Levrier, F., & Jeffrey, N. 2021, A&A, 649, L18
Régaldo-Saint Blancard, B., Hahn, C., Ho, S., et al. 2024, Physical Review D, 109, 083535
Regaldo-Saint Blancard, B., Levrier, F., Allys, E., Bellomi, E., & Boulanger, F. 2020, A&A, 642, A217
Roddy, P. J. & McEwen, J. D. 2021, IEEE Sig. Proc. Let., 28, 304
Saydjari, A. K., Portillo, S. K. N., Slepian, Z., et al. 2021, ApJ, 910, 122
Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920
Siahkoohi, A., Morel, R., Balestriero, R., et al. 2023a, arXiv preprint arXiv:2305.16189
Siahkoohi, A., Morel, R., Maarten, V., et al. 2023b, in International Conference on Machine Learning, PMLR, 31754–31772
Thorne, B., Dunkley, J., Alonso, D., & Næss, S. 2017, MNRAS, 469, 2821
Valogiannis, G. & Dvorkin, C. 2022a, Phys. Rev. D, 106, 103509
Valogiannis, G. & Dvorkin, C. 2022b, Phys. Rev. D, 105, 103534
Varshalovich, D. A., Moskalev, A. N., & Khersonskii, V. K. 1988, Quantum Theory of Angular Momentum
Zhang, S. & Mallat, S. 2021, Applied and Computational Harmonic Analysis, 53, 199

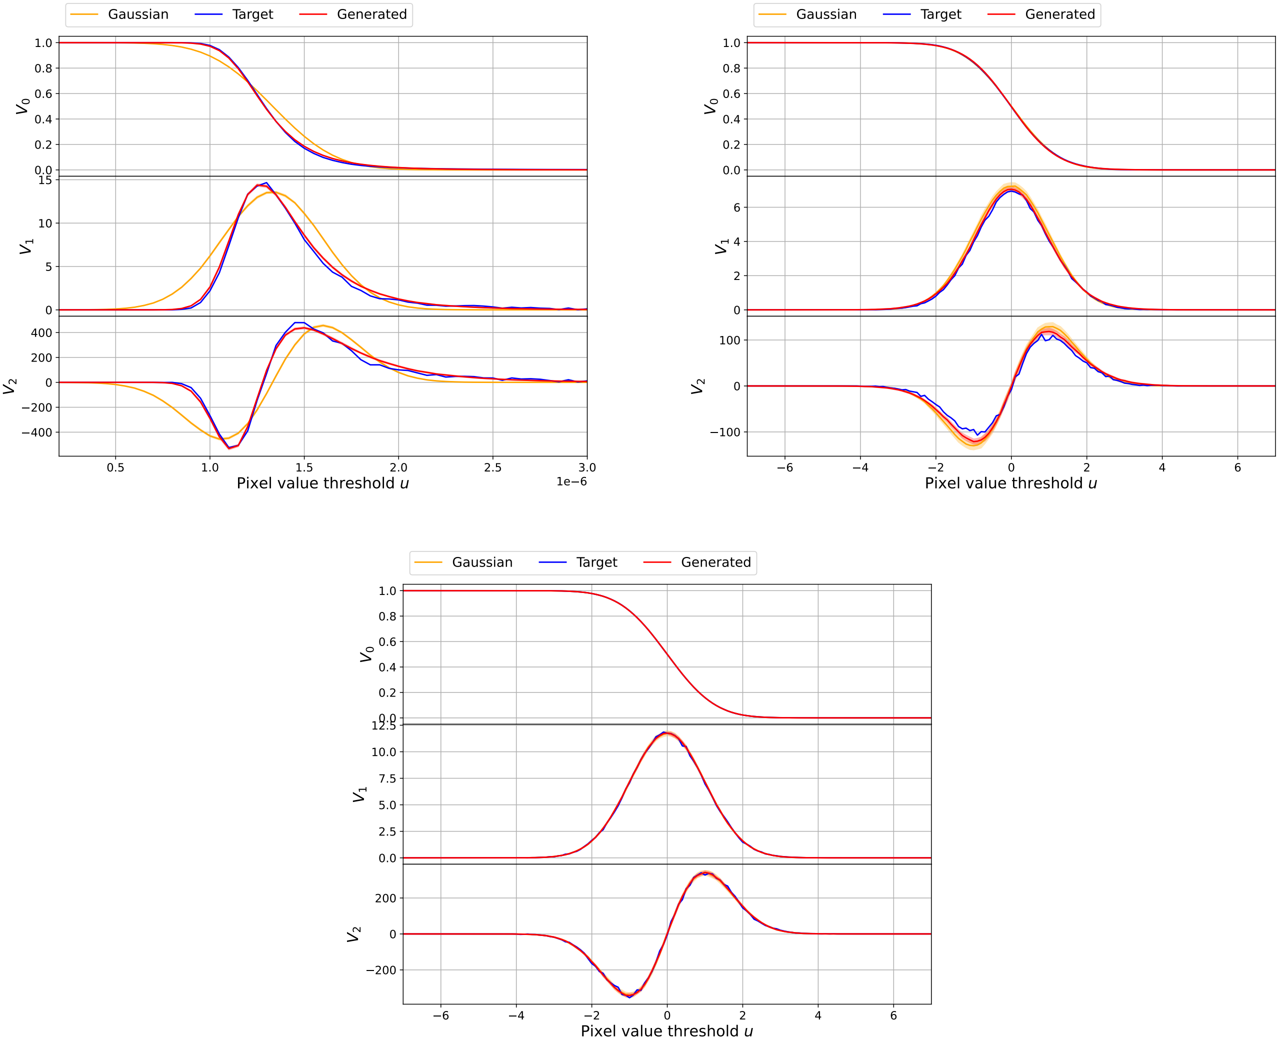

Fig. D.1. The three Minkowski functionals V_0, V_1 and V_2 for the tSZ (upper left), Venus (upper right) and CMB (bottom) fields. Blue is the target, red the generated field and yellow the Gaussian fields. For the generated field and the Gaussian realisations, we plot the mean (solid) and the standard deviation (shadow envelope) computed over 50 realisations.